Journal of Machine Learning Research 17 (2016) 1-49 Submitted 5/14; Revised 1/16; Published 5/16

SPSD Matrix Approximation vis Column Selection: Theories, Algorithms, and Extensions

Shusen Wang WSSATZJU@GMAIL.COM
Department of Statistics
University of California at Berkeley
Berkeley, CA 94720

Luo Luo RICKY@SJTU.EDU.CN
Zhihua Zhang* ZHIHUA@SJTU.EDU.CN
Department of Computer Science and Engineering
Shanghai Jiao Tong University
800 Dong Chuan Road, Shanghai, China 200240

Editor: Inderjit Dhillon

arXiv:1406.5675v6 [cs.LG] 20 May 2016

Abstract

Symmetric positive semidefinite (SPSD) matrix approximation is an important problem with applications in kernel methods. However, existing SPSD matrix approximation methods such as the Nyström method only have weak error bounds. In this paper we conduct in-depth studies of an SPSD matrix approximation model and establish strong relative-error bounds. We call it the prototype model for it has more efficient and effective extensions, and some of its extensions have high scalability. Though the prototype model itself is not suitable for large-scale data, it is still useful to study its properties, on which the analysis of its extensions relies.

This paper offers novel theoretical analysis, efficient algorithms, and a highly accurate extension. First, we establish a lower error bound for the prototype model and improve the error bound of an existing column selection algorithm to match the lower bound. In this way, we obtain the first optimal column selection algorithm for the prototype model. We also prove that the prototype model is exact under certain conditions. Second, we develop a simple column selection algorithm with a provable error bound. Third, we propose a so-called spectral shifting model to make the approximation more accurate when the eigenvalues of the matrix decay slowly, and the improvement is theoretically quantified. The spectral shifting method can also be applied to improve other SPSD matrix approximation models.

Keywords: Matrix approximation, matrix factorization, kernel methods, the Nyström method, spectral shifting

1. Introduction

The kernel methods are important tools in machine learning, computer vision, and data mining (Schölkopf and Smola, 2002; Shawe-Taylor and Cristianini, 2004). However, for two reasons, most kernel methods have scalability difficulties. First, given n data points of d

*. Corresponding author.

©2016 Shusen Wang, Luo Luo, and Zhihua Zhang.

dimension, generally we need $\mathcal{O}(n^2 d)$ time to form the $n \times n$ kernel matrix $\mathbf{K}$. Second, most kernel methods require expensive matrix computations. For example, Gaussian process regression and classification require inverting some $n \times n$ matrices which costs $\mathcal{O}(n^3)$ time and $\mathcal{O}(n^2)$ memory; kernel PCA and spectral clustering perform the truncated eigenvalue decomposition which takes $\tilde{\mathcal{O}}(n^2 k)$ time[1] and $\mathcal{O}(n^2)$ memory, where k is the target rank of the decomposition.

Besides high time complexities, these matrix operations also have high memory cost and are difficult to implement in distributed computing facilities. The matrix decomposition and (pseudo) inverse operations are generally solved by numerical iterative algorithms, which go many passes through the matrix until convergence. Thus, the whole matrix had better been placed in main memory, otherwise in each iteration there would be a swap between memory and disk, which incurs high I/O costs and can be more expensive than CPU time. Unless the algorithm is pass-efficient, that is, it goes constant passes through the data matrix, the main memory should be at least the size of the data matrix. For two reasons, such iterative algorithms are expensive even if they are performed in distributed computing facilities such as MapReduce. First, the memory cost is too expensive for each individual machine to stand. Second, communication and synchronization must be performed in each iteration of the numerical algorithms, so the cost of each iteration is high.

Many matrix approximation methods have been proposed to make kernel machines scalable. Among them the Nyström method (Nyström, 1930; Williams and Seeger, 2001) and random features (Rahimi and Recht, 2008) are the most efficient and widely applied. However, only weak results are known (Drineas and Mahoney, 2005; Gittens and Mahoney, 2013; Lopez-Paz et al., 2014). Yang et al. (2012) showed that the Nyström method is likely a better choice than random features, both theoretically and empirically. However, even the Nyström method cannot attain high accuracy. The lower bound in (Wang and Zhang, 2013) indicates that the Nyström method costs at least $\Omega(n^2 k/\epsilon)$ time and $\Omega(n^{1.5} k^{0.5} \epsilon^{-0.5})$ memory to attain $1+\epsilon$ Frobenius norm error bound relative to the best rank k approximation.

In this paper we investigate a more accurate low-rank approximation model proposed by Halko et al. (2011); Wang and Zhang (2013), which we refer to as the prototype model. For any symmetric positive semidefinite (SPSD) matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$, the prototype model first draws a random matrix $\mathbf{P} \in \mathbb{R}^{n \times c}$ and forms a sketch $\mathbf{C} = \mathbf{K}\mathbf{P}$, and then computes the intersection matrix

$$\mathbf{U}^\star \;=\; \operatorname*{argmin}_{\mathbf{U}} \|\mathbf{K} - \mathbf{C}\mathbf{U}\mathbf{C}^T\|_F^2 \;=\; \mathbf{C}^\dagger \mathbf{K} (\mathbf{C}^\dagger)^T \;\in\; \mathbb{R}^{c \times c}. \tag{1}$$

Finally, the model approximates $\mathbf{K}$ by $\mathbf{C}\mathbf{U}^\star\mathbf{C}^T$. With this low-rank approximation at hand, it takes time $\mathcal{O}(nc^2)$ to compute the approximate matrix inversion and eigenvalue decomposition. In the following we discuss how to form $\mathbf{C}$ and compute $\mathbf{U}^\star$.

Column Selection vs. Random Projection. Although the sketch $\mathbf{C} = \mathbf{K}\mathbf{P}$ can be formed by either random projection or column selection, when applied to the kernel methods, column selection is preferable to random projection. As aforementioned, suppose we are given n data points of d dimension. It takes time $\mathcal{O}(n^2 d)$ to compute the whole of the kernel matrix $\mathbf{K}$, which is prohibitive when n is in million scale. Unfortunately, whatever

1. The $\tilde{\mathcal{O}}$ notation hides the logarithm terms and the data-dependent spectral gap parameter.

existing random projection technique is employed to form the sketch $\mathbf{C}$, every entry of $\mathbf{K}$ must be visited. In contrast, by applying data independent column selection algorithms such as uniform sampling, we can form $\mathbf{C}$ by observing only $\mathcal{O}(nc)$ entries of $\mathbf{K}$. At present all the existing column selection algorithms, including our proposed uniform+adaptive^2 algorithm, cannot avoid observing the whole of $\mathbf{K}$ while keeping constant-factor bound. Nevertheless, we conjecture that our uniform+adaptive^2 algorithm can be adapted to satisfy these two properties simultaneously (see Section 5.4 for discussions in detail).

The Intersection Matrix. With the sketch $\mathbf{C}$ at hand, it remains to compute the intersection matrix. The most straightforward way is (1), which minimizes the Frobenius norm approximation error. However, this approach has two drawbacks. First, it again requires the full observation of $\mathbf{K}$. Second, the matrix product $\mathbf{C}^\dagger\mathbf{K}$ costs $\mathcal{O}(n^2c)$ time. The prototype model is therefore time-inefficient. Fortunately, Wang et al. (2015) recently overcame the two drawbacks by solving (1) approximately rather than optimally. Wang et al. (2015) obtained the approximate intersection matrix $\tilde{\mathbf{U}}$ in $\mathcal{O}(nc^3/\epsilon)$ time while keeping

$$\|\mathbf{K} - \mathbf{C}\tilde{\mathbf{U}}\mathbf{C}^T\|_F^2 \;\le\; (1+\epsilon)\min_{\mathbf{U}} \|\mathbf{K} - \mathbf{C}\mathbf{U}\mathbf{C}^T\|_F^2 \tag{2}$$

with high probability.

With the more efficient solution, why is it useful to study the exact solution to the prototype model (1)? On the one hand, from (2) we can see that the quality of the approximation depends on the prototype model, thus improvement of the prototype model directly applies to the more efficient model. On the other hand, for medium-scale problems where $\mathbf{K}$ does not fit in memory, the prototype model can produce high quality approximation with reasonable time expense. The experiment on kernel PCA in Section 6.3 shows that the prototype model is far more accurate than the Nyström method. For the above two reasons, we believe the study of the prototype model is useful.

1.1 Contributions

Our contributions mainly include three aspects: theoretical analysis, column selection algorithms, and extensions. They are summarized as follows.

1.1.1 Contributions: Theories

Kumar et al. (2009); Talwalkar and Rostamizadeh (2010) previously showed that the Nyström method is exact when the original kernel matrix is low-rank. In Section 4.1 we show that the prototype model exactly recovers the original SPSD matrix under the same conditions.

The prototype model with the near-optimal+adaptive column sampling algorithm satisfies $1+\epsilon$ relative-error bound when $c = \mathcal{O}(k/\epsilon^2)$ (Wang and Zhang, 2013). It was unknown whether this upper bound is optimal. In Section 4.2 we establish a lower error bound for the prototype model. We show that at least $2k/\epsilon$ columns must be chosen to attain $1+\epsilon$ bound. In Theorem 3 we improve the upper error bound of the near-optimal+adaptive algorithm to $\mathcal{O}(k/\epsilon)$, which matches the lower bound up to a constant factor.

1.1.2 Contributions: Algorithms

In Section 5 we devise a simple column selection algorithm which we call *the uniform+adaptive^2 algorithm*. The uniform+adaptive^2 algorithm is more efficiently and more easily imple-

mented than the near-optimal+adaptive algorithm of Wang and Zhang (2013), yet its error bound is comparable with the near-optimal+adaptive algorithm. It is worth mentioning that our uniform+adaptive^2 algorithm is the adaptive-full algorithm of (Figure 3, Kumar et al., 2012) with two rounds of adaptive sampling, and thus our results theoretically justify the adaptive-full algorithm.

1.1.3 Contributions: Extension

When the spectrum of a matrix decays slowly (that is, the $c+1$ to n largest eigenvalues are not small enough), all of the low-rank approximations are far from the original kernel matrix. Inspired by Zhang (2014), we propose a new method called *spectral shifting (SS)* to make the approximation still effective even when the spectrum decays slowly. Unlike the low-rank approximation $\mathbf{K} \approx \mathbf{C}\mathbf{U}\mathbf{C}^T$, the spectral shifting model approximates $\mathbf{K}$ by $\mathbf{K} \approx \bar{\mathbf{C}}\mathbf{U}^{\mathrm{ss}}\bar{\mathbf{C}}^T + \delta^{\mathrm{ss}}\mathbf{I}_n$, where $\mathbf{C}, \bar{\mathbf{C}} \in \mathbb{R}^{n\times c}$, $\mathbf{U}, \mathbf{U}^{\mathrm{ss}} \in \mathbb{R}^{c\times c}$, and $\delta^{\mathrm{ss}} \geq 0$. When the spectrum of $\mathbf{K}$ decays slowly, the term $\delta^{\mathrm{ss}}\mathbf{I}_n$ helps to improve the approximation accuracy. In Section 7 we describe the spectral shifting method in detail.

We highlight that the spectral shifting method can naturally apply to improve other kernel approximation models such as the memory efficient kernel approximation (MEKA) model (Si et al., 2014). Experiments demonstrate that MEKA can be significantly improved by spectral shifting.

1.2 Paper Organization

The remainder of this paper is organized as follows. In Section 2 we define the notation. In Section 3 we introduce the motivations of SPSD matrix approximation and define the SPSD matrix approximation models. Then we present our work—theories, algorithms, and extension—respectively in Sections 4, 5, and 7. In Section 6 we conduct experiments to compare among the column sampling algorithms. In Section 8 we empirically evaluate the proposed spectral shifting model. All the proofs are deferred to the appendix.

2. Notation

Let $[n] = \{1, \ldots, n\}$, and $\mathbf{I}_n$ be the $n\times n$ identity matrix. For an $m\times n$ matrix $\mathbf{A} = [a_{ij}]$, we let $\mathbf{a}_{i:}$ be its i-th row, $\mathbf{a}_{:i}$ be its i-th column, and use $\mathbf{a}_i$ to denote either row or column when there is no ambiguity. Let $\mathbf{A}_1 \oplus \mathbf{A}_2 \oplus \cdots \oplus \mathbf{A}_q$ be the block diagonal matrix whose the i-th diagonal block is $\mathbf{A}_i$. Let $\|\mathbf{A}\|_F = (\sum_{i,j} a_{ij}^2)^{1/2}$ be the Frobenius norm and $\|\mathbf{A}\|_2 = \max_{\mathbf{x}\neq\mathbf{0}} \|\mathbf{A}\mathbf{x}\|_2/\|\mathbf{x}\|_2$ be the spectral norm.

Letting $\rho = \mathrm{rank}(\mathbf{A})$, we write the condensed singular value decomposition (SVD) of $\mathbf{A}$ as $\mathbf{A} = \mathbf{U}_{\mathbf{A}}\mathbf{\Sigma}_{\mathbf{A}}\mathbf{V}_{\mathbf{A}}^T$, where the (i,i)-th entry of $\mathbf{\Sigma}_{\mathbf{A}} \in \mathbb{R}^{\rho\times\rho}$ is the i-th largest singular value of $\mathbf{A}$ (denoted $\sigma_i(\mathbf{A})$). Unless otherwise specified, in this paper "SVD" means the condensed SVD. We also let $\mathbf{U}_{\mathbf{A},k}$ and $\mathbf{V}_{\mathbf{A},k}$ be the first k ($< \rho$) columns of $\mathbf{U}_{\mathbf{A}}$ and $\mathbf{V}_{\mathbf{A}}$, respectively, and $\mathbf{\Sigma}_{\mathbf{A},k}$ be the $k\times k$ top sub-block of $\mathbf{\Sigma}_{\mathbf{A}}$. Then the $m\times n$ matrix $\mathbf{A}_k = \mathbf{U}_{\mathbf{A},k}\mathbf{\Sigma}_{\mathbf{A},k}\mathbf{V}_{\mathbf{A},k}^T$ is the "closest" rank-k approximation to $\mathbf{A}$.

If $\mathbf{A}$ is normal, we let $\mathbf{A} = \mathbf{U}_{\mathbf{A}}\mathbf{\Lambda}_{\mathbf{A}}\mathbf{U}_{\mathbf{A}}^T$ be the eigenvalue decomposition, and denote the i-th diagonal entry of $\mathbf{\Lambda}_{\mathbf{A}}$ by $\lambda_i(\mathbf{A})$, where $|\lambda_1(\mathbf{A})| \geq \cdots \geq |\lambda_n(\mathbf{A})|$. When $\mathbf{A}$ is SPSD, the SVD and the eigenvalue decomposition of $\mathbf{A}$ are identical.

Table 1: Comparisons among the matrix approximation models in Section 3.2 and Section 3.3. Here "#Entries" denotes the number of entries of $\mathbf{K}$ required to observe. The costs of column selection is not counted; they are listed separately in Table 2.

	Time	Memory	#Entries	Theory
Prototype	$\mathcal{O}(n^2c)$	$\mathcal{O}(nc)$	n^2	$1+\epsilon$ relative-error
Faster	$\mathcal{O}(nc^3/\epsilon)$	$\mathcal{O}(nc)$	nc^2/ϵ	$1+\epsilon$ relative-error
Nyström	$\mathcal{O}(nc^2)$	$\mathcal{O}(nc)$	nc	weak
SS	the same to "prototype"			stronger than "prototype"
Faster SS	the same to "faster"			unknown

Based on SVD, the *matrix coherence* of the columns of $\mathbf{A}$ relative to the best rank-k approximation is defined as $\mu_k = \frac{n}{k}\max_j \left\|(\mathbf{V}_{\mathbf{A},k})_{j:}\right\|_2^2$. Let $\mathbf{A}^\dagger = \mathbf{V}_\mathbf{A}\mathbf{\Sigma}_\mathbf{A}^{-1}\mathbf{U}_\mathbf{A}^T$ be the *Moore-Penrose inverse* of $\mathbf{A}$. When $\mathbf{A}$ is nonsingular, the Moore-Penrose inverse is identical to the matrix inverse. Given another $n \times c$ matrix $\mathbf{C}$, we define $\mathcal{P}_\mathbf{C}(\mathbf{A}) = \mathbf{C}\mathbf{C}^\dagger\mathbf{A}$ as the projection of $\mathbf{A}$ onto the column space of $\mathbf{C}$ and $\mathcal{P}_{\mathbf{C},k}(\mathbf{A}) = \mathbf{C}\cdot\operatorname{argmin}_{\operatorname{rank}(\mathbf{X})\le k}\|\mathbf{A}-\mathbf{C}\mathbf{X}\|_F$ as the rank restricted projection. It is obvious that $\|\mathbf{A}-\mathcal{P}_\mathbf{C}(\mathbf{A})\|_F \le \|\mathbf{A}-\mathcal{P}_{\mathbf{C},k}(\mathbf{A})\|_F$.

3. SPSD Matrix Approximation Models

In Section 3.1 we provide motivating examples to show why SPSD matrix approximation is useful. In Section 3.2 we formally describe low-rank approximation models. In Section 3.3 we describe the spectral shifting model. In Table 1 we compare the matrix approximation models defined in Section 3.2 and Section 3.3.

3.1 Motivations

Let $\mathbf{K}$ be an $n\times n$ kernel matrix. Many kernel methods require the eigenvalue decomposition of $\mathbf{K}$ or solving certain linear systems involving $\mathbf{K}$.

- Spectral clustering, kernel PCA, and manifold learning need to perform the rank k eigenvalue decomposition which costs $\tilde{\mathcal{O}}(n^2k)$ time and $\mathcal{O}(n^2)$ memory.
- Gaussian process regression and classification both require solving this kind of linear systems:
$$(\mathbf{K}+\alpha\mathbf{I}_n)\mathbf{b} = \mathbf{y}, \tag{3}$$
whose solution is $\mathbf{b}^\star = (\mathbf{K}+\alpha\mathbf{I}_n)^{-1}\mathbf{y}$. Here α is a constant. This costs $\mathcal{O}(n^3)$ time and $\mathcal{O}(n^2)$ memory.

Fortunately, if we can efficiently find an approximation in the form
$$\tilde{\mathbf{K}} \;=\; \mathbf{L}\mathbf{L}^T + \delta\mathbf{I}_n \;\approx\; \mathbf{K},$$
where $\delta \ge 0$ and $\mathbf{L} \in \mathbb{R}^{n\times l}$ with $l \ll n$, then the eigenvalue decomposition and linear systems can be approximately solved in $\mathcal{O}(nl^2)$ time and $\mathcal{O}(nl)$ space in the following way.

Algorithm 1 Computing the Prototype Model in $\mathcal{O}(nc+nd)$ Memory.

Input: data points $\mathbf{x}_1, \cdots, \mathbf{x}_n \in \mathbb{R}^d$, kernel function $\kappa(\cdot,\cdot)$.
Compute $\mathbf{C}$ and $\mathbf{C}^\dagger$; *// In $\mathcal{O}(nc+nd)$ memory and $\mathcal{O}(ncd+nc^2)$ time*
Form a $c \times n$ all-zero matrix $\mathbf{D}$; *// In $\mathcal{O}(nc)$ memory and $\mathcal{O}(nc)$ time*
for $j = 1$ to n **do**
    Form the j-th column of $\mathbf{K}$ by $\mathbf{k}_j = [\kappa(\mathbf{x}_1, \mathbf{x}_j), \cdots, \kappa(\mathbf{x}_n, \mathbf{x}_j)]^T$;
    Compute the j-th column of $\mathbf{D}$ by $\mathbf{d}_j = \mathbf{C}^\dagger \mathbf{k}_j$;
    Delete $\mathbf{k}_j$ from memory;
end for
// Now the matrix $\mathbf{D}$ is $\mathbf{C}^\dagger\mathbf{K}$
// The loop totaly costs $\mathcal{O}(nc+nd)$ memory and $\mathcal{O}(n^2d+n^2c)$ time
Compute $\mathbf{U} = \mathbf{D}(\mathbf{C}^\dagger)^T$; *// In $\mathcal{O}(nc)$ memory and $\mathcal{O}(nc^2)$ time*
return $\mathbf{C}$ and $\mathbf{U}$ $(= \mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T)$.

- Approximate Eigenvalue Decomposition. Let $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ be the SVD and $\mathbf{U}_\perp$ be the orthogonal complement of $\mathbf{U}$. Then the full eigenvalue decomposition of $\tilde{\mathbf{K}}$ is

$$\tilde{\mathbf{K}} \;=\; \mathbf{U}(\mathbf{\Sigma}^2 + \delta\mathbf{I}_l)\mathbf{U}^T + \mathbf{U}_\perp(\delta\mathbf{I}_{n-l})\mathbf{U}_\perp^T.$$

- Approximately Solving the Linear Systems. Here we use a more general form: $\tilde{\mathbf{K}} = \mathbf{L}\mathbf{L}^T + \mathbf{\Delta}$, where $\mathbf{\Delta}$ is a diagonal matrix with positive diagonal entries. Then

$$\begin{aligned}\mathbf{b}^\star \;&=\; (\mathbf{K}+\alpha\mathbf{I}_n)^{-1}\mathbf{y} \;\approx\; (\mathbf{L}\mathbf{L}^T + \mathbf{\Delta} + \alpha\mathbf{I}_n)^{-1}\mathbf{y} \;=\; (\mathbf{L}\mathbf{L}^T + \mathbf{\Delta}')^{-1}\mathbf{y} \\ &=\; \mathbf{\Delta}'^{-1}\mathbf{y} - \underbrace{\mathbf{\Delta}'^{-1}\mathbf{L}}_{n\times l}\underbrace{(\mathbf{I}_l + \mathbf{L}^T\mathbf{\Delta}'^{-1}\mathbf{L})^{-1}}_{l\times l}\underbrace{\mathbf{L}^T\mathbf{\Delta}'^{-1}}_{l\times n}\mathbf{y}.\end{aligned}$$

  Here the second equality is obtained by letting $\mathbf{\Delta}' = \mathbf{\Delta} + \alpha\mathbf{I}_n$, and the third equality follows by the Sherman-Morrison-Woodbury matrix identity.

The remaining problem is to find such matrix approximation efficiently while keeping $\tilde{\mathbf{K}}$ close to $\mathbf{K}$.

3.2 Low-Rank Matrix Approximation Models

We first recall the prototype model introduced previously and then discuss its approximate solutions. In fact, the famous Nyström method (Nyström, 1930; Williams and Seeger, 2001) is an approximation to the prototype model. Throughout this paper, we let $\mathbf{P} \in \mathbb{R}^{n\times c}$ be random projection or column selection matrix and $\mathbf{C} = \mathbf{K}\mathbf{P}$ be a sketch of $\mathbf{K}$. The only difference among the discussed models is their intersection matrices.

The Prototype Model. Suppose we have $\mathbf{C} \in \mathbb{R}^{n\times c}$ at hand. It remains to find an intersection matrix $\mathbf{U} \in \mathbb{R}^{c\times c}$. Since our objective is to make $\mathbf{C}\mathbf{U}\mathbf{C}^T$ close to $\mathbf{K}$, it is straightforward to optimize their difference. The prototype model computes the intersection matrix by

$$\mathbf{U}^\star \;=\; \operatorname*{argmin}_{\mathbf{U}} \|\mathbf{K} - \mathbf{C}\mathbf{U}\mathbf{C}^T\|_F^2 \;=\; \mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger \;\in\; \mathbb{R}^{c\times c}. \tag{4}$$

With $\mathbf{C}$ at hand, the prototype model still needs one pass through the data, and it costs $\mathcal{O}(n^2c)$ time. When applied to kernel methods, the memory cost is $\mathcal{O}(nc+nd)$ (see Algorithm 1), where n is the number of data points and d is the dimension. The prototype model has the same time complexity as the exact rank k eigenvalue decomposition, but it is more memory-efficient and pass-efficient.

Halko et al. (2011) showed that when $\mathbf{P}$ is a standard Gaussian matrix and $c=\mathcal{O}(k/\epsilon)$, the prototype model attains $2+\epsilon$ error relative to $\|\mathbf{K}-\mathbf{K}_k\|_F^2$. Wang and Zhang (2013) showed that when $\mathbf{C}$ contains $c=\mathcal{O}(k/\epsilon^2)$ columns selected by the near-optimal+adaptive sampling algorithm, the prototype model attains $1+\epsilon$ relative error. In Section 5.2 we improve the result to $c=\mathcal{O}(k/\epsilon)$, which is near optimal.

Faster SPSD matrix Approximation Model. Wang et al. (2015) noticed that (4) is a strongly over-determined linear system, and thus proposed to solve (4) by randomized approximations. They proposed to sample $s=\mathcal{O}(c\sqrt{n/\epsilon}) \ll n$ columns according to the row leverage scores of $\mathbf{C}$, which costs $\mathcal{O}(nc^2)$ time. Let $\mathbf{S}\in\mathbb{R}^{n\times s}$ be the corresponding column selection matrix. They proposed the faster SPSD matrix approximation model which computes the intersection matrix by

$$\tilde{\mathbf{U}} \;=\; \operatorname*{argmin}_{\mathbf{U}} \left\|\mathbf{S}^T(\mathbf{K}-\mathbf{C}\mathbf{U}\mathbf{C}^T)\mathbf{S}\right\|_F^2 \;=\; \underbrace{(\mathbf{S}^T\mathbf{C})^\dagger}_{c\times s}\underbrace{(\mathbf{S}^T\mathbf{K}\mathbf{S})}_{s\times s}\underbrace{(\mathbf{C}^T\mathbf{S})^\dagger}_{s\times c} \;\in\; \mathbb{R}^{c\times c}. \tag{5}$$

The faster model visits only $s^2=\mathcal{O}(nc^2/\epsilon)=o(n^2)$ entries of $\mathbf{K}$, and the time complexity is $\mathcal{O}(nc^2+s^2c)=\mathcal{O}(nc^3/\epsilon)$. The following error bound is satisfied with high probability:

$$\|\mathbf{K}-\mathbf{C}\tilde{\mathbf{U}}\mathbf{C}^T\|_F^2 \;\le\; (1+\epsilon)\min_{\mathbf{U}}\|\mathbf{K}-\mathbf{C}\mathbf{U}\mathbf{C}^T\|_F^2.$$

This implies that if $\mathbf{C}$ is such a high quality sketch that the prototype model satisfies $1+\epsilon$ relative-error bound, then the faster SPSD matrix approximation model also satisfies $1+\epsilon$ relative-error bound.

The Nyström Method. The Nyström method is a special case of the faster SPSD matrix approximation model, and therefore it is also an approximate solution to (4). If we let the two column selection matrices $\mathbf{S}$ and $\mathbf{P}$ be the same, then (5) becomes

$$\tilde{\mathbf{U}} \;=\; \operatorname*{argmin}_{\mathbf{U}} \left\|\mathbf{P}^T(\mathbf{K}-\mathbf{C}\mathbf{U}\mathbf{C}^T)\mathbf{P}\right\|_F^2 \;=\; \underbrace{(\mathbf{P}^T\mathbf{K}\mathbf{P})^\dagger}_{c\times c}\underbrace{(\mathbf{P}^T\mathbf{K}\mathbf{P})}_{c\times c}\underbrace{(\mathbf{P}^T\mathbf{K}\mathbf{P})^\dagger}_{c\times c} \;=\; \underbrace{(\mathbf{P}^T\mathbf{K}\mathbf{P})^\dagger}_{c\times c}.$$

The matrix $(\mathbf{P}^T\mathbf{K}\mathbf{P})^\dagger$ is exactly the intersection matrix of the Nyström method. The Nyström method costs only $\mathcal{O}(nc^2)$ time, and it can be applied to million-scale problems (Talwalkar et al., 2013). However, its accuracy is low. Much work in the literature has analyzed the error bound of the Nyström method, but only weak results are known (Drineas and Mahoney, 2005; Shawe-Taylor et al., 2005; Kumar et al., 2012; Jin et al., 2013; Gittens and Mahoney, 2013). Wang and Zhang (2013) even showed that the Nyström method cannot attain $1+\epsilon$ relative-error bound unless $c\ge\Omega(\sqrt{nk/\epsilon})$. Equivalently, to attain $1+\epsilon$ bound, the Nyström would take $\Omega(n^2k/\epsilon)$ time and $\Omega(n^{1.5}k^{0.5}\epsilon^{-0.5})$ memory.

3.3 Spectral Shifting Models

We propose a more accurate SPSD matrix approximation method called the spectral shifting model. Here we briefly describe the model and its fast solution. The theoretical analysis is left to Section 7.

The Spectral Shifting (SS) Model. As before, we let $\mathbf{P} \in \mathbb{R}^{n\times c}$ be a column selection matrix and $\bar{\mathbf{C}} = \bar{\mathbf{K}}\mathbf{P}$, where $\bar{\mathbf{K}} = \mathbf{K}$ or $\bar{\mathbf{K}} = \mathbf{K} - \bar{\delta}\mathbf{I}_n$ for some parameter $\bar{\delta} \geq 0$. We approximate $\mathbf{K}$ by $\bar{\mathbf{C}}\mathbf{U}^{\text{ss}}\bar{\mathbf{C}}^T + \delta^{\text{ss}}\mathbf{I}_n$, where

$$\big(\mathbf{U}^{\text{ss}}, \delta^{\text{ss}}\big) \;=\; \operatorname*{argmin}_{\mathbf{U},\delta} \big\|\mathbf{K} - \bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T - \delta\mathbf{I}_n\big\|_F^2. \tag{6}$$

This optimization problem has closed-form solution (see Theorem 6)

$$\begin{aligned}\delta^{\text{ss}} &= \frac{1}{n - \operatorname{rank}(\bar{\mathbf{C}})}\Big(\operatorname{tr}(\mathbf{K}) - \operatorname{tr}\big(\bar{\mathbf{C}}^\dagger\mathbf{K}\bar{\mathbf{C}}\big)\Big),\\ \mathbf{U}^{\text{ss}} &= \bar{\mathbf{C}}^\dagger\mathbf{K}(\bar{\mathbf{C}}^\dagger)^T - \delta^{\text{ss}}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger,\end{aligned}$$

which can be computed in $\mathcal{O}(n^2c)$ time and $\mathcal{O}(nc)$ memory. Later we will show that the SS model is more accurate than the prototype model.

Faster Spectral Shifting Model. The same idea of Wang et al. (2015) also applies to the SS model (14). Specifically, we can draw another column selection matrix $\mathbf{S} \in \mathbb{R}^{n\times s}$ and solve

$$\begin{aligned}\big(\tilde{\mathbf{U}}^{\text{ss}}, \tilde{\delta}^{\text{ss}}\big) &= \operatorname*{argmin}_{\mathbf{U},\delta} \big\|\mathbf{S}^T\big(\mathbf{K} - \bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T - \delta\mathbf{I}_n\big)\mathbf{S}\big\|_F^2\\ &= \operatorname*{argmin}_{\mathbf{U},\delta} \big\|\mathbf{S}^T\mathbf{K}\mathbf{S} - (\mathbf{S}^T\bar{\mathbf{C}})\mathbf{U}(\mathbf{S}^T\bar{\mathbf{C}})^T - \delta\mathbf{I}_s\big\|_F^2.\end{aligned}$$

Similarly, it has closed-form solution

$$\begin{aligned}\tilde{\delta}^{\text{ss}} &= \frac{1}{s - \operatorname{rank}(\mathbf{S}^T\bar{\mathbf{C}})}\Big[\operatorname{tr}\big(\mathbf{S}^T\mathbf{K}\mathbf{S}\big) - \operatorname{tr}\Big((\mathbf{S}^T\bar{\mathbf{C}})^\dagger(\mathbf{S}^T\mathbf{K}\mathbf{S})(\mathbf{S}^T\bar{\mathbf{C}})\Big)\Big],\\ \tilde{\mathbf{U}}^{\text{ss}} &= (\mathbf{S}^T\bar{\mathbf{C}})^\dagger(\mathbf{S}^T\mathbf{K}\mathbf{S})(\bar{\mathbf{C}}^T\mathbf{S})^\dagger - \tilde{\delta}^{\text{ss}}(\bar{\mathbf{C}}^T\mathbf{S}\mathbf{S}^T\bar{\mathbf{C}})^\dagger.\end{aligned}$$

In this way, the time cost is merely $\mathcal{O}(s^2c)$. However, the theoretical properties of this model are yet unknown. We do not conduct theoretical or empirical study of this model; we leave it as a future work,

4. Theories

In Section 4.1 we show that the prototype model is exact when $\mathbf{K}$ is low-rank. In Section 4.2 we provide a lower error bound of the prototype model.

4.1 Theoretical Justifications

Let $\mathbf{P}$ be a column selection matrix, $\mathbf{C} = \mathbf{K}\mathbf{P}$ be a sketch, and $\mathbf{W} = \mathbf{P}^T\mathbf{K}\mathbf{P}$ be the corresponding submatrix. Kumar et al. (2009); Talwalkar and Rostamizadeh (2010) showed that the Nyström method is exact when $\operatorname{rank}(\mathbf{W}) = \operatorname{rank}(\mathbf{K})$. We present a similar result in Theorem 1.

Theorem 1 *The following three statements are equivalent:* (i) $\operatorname{rank}(\mathbf{W}) = \operatorname{rank}(\mathbf{K})$, (ii) $\mathbf{K} = \mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T$, (iii) $\mathbf{K} = \mathbf{C}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T\mathbf{C}^T$.

Theorem 1 implies that the prototype model and the Nyström method are equivalent when $\operatorname{rank}(\mathbf{W}) = \operatorname{rank}(\mathbf{K})$; that is, the kernel matrix $\mathbf{K}$ is low rank. However, it holds in general that $\operatorname{rank}(\mathbf{K}) \gg c \geq \operatorname{rank}(\mathbf{W})$, where the two models are not equivalent.

4.2 Lower Bound

Wang and Zhang (2013) showed that with $c = \mathcal{O}(k/\epsilon^2)$ columns chosen by the near-optimal+adaptive sampling algorithm, the prototype model satisfies $1+\epsilon$ relative-error bound. We establish a lower error bound in Theorem 2, which shows that at least $c \geq 2k\epsilon^{-1}$ columns must be chosen to attain the $1+\epsilon$ bound. This indicates there exists a gap between the upper bound in (Wang and Zhang, 2013) and our lower bound, and thus there is room of improvement. The proof of Theorem 2 is left to Appendix B.

Theorem 2 (Lower Bound of the Prototype Model) *Whatever column sampling algorithm is used, there exists an $n \times n$ SPSD matrix $\mathbf{K}$ such that the error incurred by the prototype model obeys:*

$$\left\|\mathbf{K} - \mathbf{C}\mathbf{U}^\star\mathbf{C}^T\right\|_F^2 \quad \geq \quad \frac{n-c}{n-k}\left(1+\frac{2k}{c}\right)\|\mathbf{K}-\mathbf{K}_k\|_F^2.$$

Here k is an arbitrary target rank, c is the number of selected columns, and $\mathbf{U}^\star = \mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T$.

5. Column Sampling Algorithms

In Section 5.1 we introduce the column sampling algorithms in the literature. In Section 5.2 we improve the bound of the near-optimal+adaptive sampling algorithm (Wang and Zhang, 2013), and the obtained upper bound matches the lower bound up to a constant factor. In Section 5.3 we develop a more efficient column sampling algorithm which we call the uniform+adaptive^2 algorithm. In Section 5.4 we discuss the possibility of making uniform+adaptive^2 more scalable.

5.1 Related Work

Column selection is an important matrix sketching approach that enables expensive matrix computations to be performed on much a smaller matrix. The column selection problem has been widely studied in the theoretical computer science community (Boutsidis et al., 2014; Mahoney, 2011; Guruswami and Sinop, 2012; Woodruff, 2014) and the numerical linear algebra community (Gu and Eisenstat, 1996; Stewart, 1999), and numerous algorithms have been devised and analyzed. Here we focus on some provable algorithms studied in the theoretical computer science community.

The adaptive sampling algorithm devised by Deshpande et al. (2006) (see Algorithm 2) is the most relevant to this paper. The adaptive sampling algorithm has strong error bound (Deshpande et al., 2006; Wang and Zhang, 2013; Boutsidis et al., 2014) and good empirical performance (Kumar et al., 2012). Particularly, Wang and Zhang (2013) proposed an algorithm that combines the near-optimal column sampling algorithm (Boutsidis et al., 2014) and the adaptive sampling algorithm (Deshpande et al., 2006). They showed that by selecting $c = \mathcal{O}(k\epsilon^{-2})$ columns of $\mathbf{K}$ to form $\mathbf{C}$, it holds that

$$\mathbb{E}\left\|\mathbf{K} - \mathbf{C}\left(\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T\right)\mathbf{C}^T\right\|_F \leq (1+\epsilon)\|\mathbf{K}-\mathbf{K}_k\|_F.$$

This error bound was the tightest among all the feasible algorithms for SPSD matrix approximation.

Algorithm 2 The Adaptive Sampling Algorithm.

1: **Input:** a residual matrix $\mathbf{B} \in \mathbb{R}^{n\times n}$ and number of selected columns c ($< n$).
2: Compute sampling probabilities $p_j = \|\mathbf{b}_{:j}\|_2^2 / \|\mathbf{B}\|_F^2$ for $j = 1, \cdots, n$;
3: Select c indices in c i.i.d. trials, in each trial the index j is chosen with probability p_j;
4: **return** an index set containing the indices of the selected columns.

5.2 Near Optimal Column Selection for SPSD Matrix Approximation

The error bound of near-optimal+adaptive can be improved by a factor of ϵ by exploiting the latest results of Boutsidis and Woodruff (2014). Using the same algorithm except for different c_2 (i.e. the number of columns selected by adaptive sampling), we obtain the following stronger theorem. Recall from Theorem 2 that the lower bound is $c \geq \Omega\big(2k\epsilon^{-1}(1+o(1))\big)$. Thus the near-optimal+adaptive algorithm is optimal up to a constant factor. The proof of the theorem is left to Appendix 3.

Theorem 3 (Near-Optimal+Adaptive) *Given a symmetric matrix $\mathbf{K} \in \mathbb{R}^{n\times n}$ and a target rank k, the algorithm samples totally $c = 3k\epsilon^{-1}\big(1+o(1)\big)$ columns of $\mathbf{K}$ to construct the approximation. Then*

$$\mathbb{E}\big\|\mathbf{K} - \mathbf{C}\big(\mathbf{C}^\dagger \mathbf{K} (\mathbf{C}^\dagger)^T\big)\mathbf{C}^T\big\|_F \;\leq\; (1+\epsilon)\|\mathbf{K}-\mathbf{K}_k\|_F.$$

The algorithm costs $\mathcal{O}\big(n^2c + nk^3\epsilon^{-2/3}\big)$ time and $\mathcal{O}(nc)$ memory in computing $\mathbf{C}$.

Despite its optimal error bound, the near-optimal+adaptive algorithm lacks of practicality. The implementation is complicated and difficult. Its main component—the near-optimal algorithm (Boutsidis et al., 2014)—is highly iterative and therefore not suitable for parallel computing. Every step of the near-optimal algorithm requires the full observation of $\mathbf{K}$ and there is no hope to avoid this. Thus we propose to use uniform sampling to replace the near-optimal algorithm. Although the obtained uniform+adaptive^2 algorithm also has quadratic time complexity and requires the full observation of $\mathbf{K}$, there may be some way to making it more efficient. See the discussions in Section 5.4.

5.3 The Uniform+Adaptive^2 Column Sampling Algorithm

In this paper we propose a column sampling algorithm which is efficient, effective, and very easy to be implemented. The algorithm consists of a uniform sampling step and two adaptive sampling steps, so we call it *the uniform+adaptive^2 algorithm*. The algorithm is described in Algorithm 3 and analyzed in Theorem 4. The proof is left to Appendix D.

It is worth mentioning that our uniform+adaptive^2 algorithm is a special instance of the adaptive-full algorithm of (Kumar et al., 2012, Figure 3). The adaptive-full algorithm consists of random initialization and multiple adaptive sampling steps. Using multiple adaptive sampling steps can surely reduce the approximation error. However, the update of sampling probability in each step is expensive, so we choose to do only two steps. The adaptive-full algorithm of (Kumar et al., 2012, Figure 3) is merely a heuristic scheme without theoretical guarantee; our result provides theoretical justification for the adaptive-full algorithm.

Algorithm 3 The Uniform+Adaptive^2 Algorithm.

1: **Input:** an $n \times n$ symmetric matrix $\mathbf{K}$, target rank k, error parameter $\epsilon \in (0, 1]$, matrix coherence μ.
2: **Uniform Sampling.** Uniformly sample
$$c_1 = 20\mu k \log\big(20k\big)$$
columns of $\mathbf{K}$ without replacement to construct $\mathbf{C}_1$;
3: **Adaptive Sampling.** Sample
$$c_2 = 17.5k/\epsilon$$
columns of $\mathbf{K}$ to construct $\mathbf{C}_2$ using the adaptive sampling algorithm (Algorithm 2) according to the residual $\mathbf{K} - \mathcal{P}_{\mathbf{C}_1}(\mathbf{K})$;
4: **Adaptive Sampling.** Sample
$$c_3 = 10k/\epsilon$$
columns of $\mathbf{K}$ to construct $\mathbf{C}_3$ using the adaptive sampling algorithm (Algorithm 2) according to the residual $\mathbf{K} - \mathcal{P}_{[\mathbf{C}_1,\, \mathbf{C}_2]}(\mathbf{K})$;
5: **return** $\mathbf{C} = [\mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3]$.

Theorem 4 (Uniform+Adaptive^2) *Given an $n \times n$ symmetric matrix $\mathbf{K}$ and a target rank k, let μ_k denote the matrix coherence of $\mathbf{K}$. Algorithm 3 samples totally*

$$c = \mathcal{O}\big(k\epsilon^{-1} + \mu_k k \log k\big)$$

columns of $\mathbf{K}$ to construct the approximation. The error bound

$$\big\|\mathbf{K} - \mathbf{C}\big(\mathbf{C}^\dagger \mathbf{K} (\mathbf{C}^\dagger)^T\big)\mathbf{C}^T\big\|_F \;\leq\; \big(1+\epsilon\big)\big\|\mathbf{K} - \mathbf{K}_k\big\|_F$$

holds with probability at least 0.7. The algorithm costs $\mathcal{O}(n^2 c)$ time and $\mathcal{O}(nc)$ space in computing $\mathbf{C}$.

Remark 5 *Theoretically, Algorithm 3 requires computing the matrix coherence of $\mathbf{K}$ in order to determine c_1 and c_2. However, computing the matrix coherence is as hard as computing the truncated SVD; even the fast approximation approach of Drineas et al. (2012) is not feasible here because $\mathbf{K}$ is a square matrix. The use of the matrix coherence here is merely for theoretical analysis; setting the parameter μ in Algorithm 3 to be exactly the matrix coherence does not certainly result in the highest accuracy. Empirically, the resulting approximation accuracy is not sensitive to the value of μ. Thus we suggest setting μ in Algorithm 3 as a constant (e.g. 1), rather than actually computing the matrix coherence.*

Table 2 presents comparisons between the near-optimal+adaptive algorithm and our uniform+adaptive^2 algorithm over the time cost, memory cost, number of passes through $\mathbf{K}$, the number of columns required to attain $1+\epsilon$ relative-error bound, and the hardness of implementation. Our algorithm is more time-efficient and pass-efficient than the near-optimal+adaptive algorithm, and the memory costs of the two algorithms are the same. To attain the same error bound, our algorithm needs to select $c = \mathcal{O}\big(k\epsilon^{-1} + \mu_k k \log k\big)$ columns, which is a little larger than that of the near-optimal+adaptive algorithm.

Table 2: Comparisons between the two sampling algorithms.

	Uniform+Adaptive^2	Near-Optimal+Adaptive
Time	$\mathcal{O}(n^2c)$	$\mathcal{O}(n^2c + nk^3\epsilon^{-2/3})$
Memory	$\mathcal{O}(nc)$	$\mathcal{O}(nc)$
#Passes	2	4
#Columns	$\mathcal{O}(k\epsilon^{-1} + \mu_k k \log k)$	$\mathcal{O}(k\epsilon^{-1})$
Implement	Easy to implement	Hard to implement

Algorithm 4 The Incomplete Uniform+Adaptive^2 Algorithm.

1: **Input:** part of an $n \times n$ symmetric matrix $\mathbf{K}$.
2: **Uniform Sampling.** Uniformly sample c_1 columns of $\mathbf{K}$ without replacement to construct $\mathbf{C}_1$;
3: **Adaptive Sampling.** Uniformly sample $o(n)$ columns of $\mathbf{K}$ to form $\mathbf{K}'$; then sample c_2 columns of $\mathbf{K}'$ to construct $\mathbf{C}_2$ using the adaptive sampling algorithm (Algorithm 2) according to the residual $\mathbf{K}' - \mathcal{P}_{\mathbf{C}_1}(\mathbf{K}')$;
4: **Adaptive Sampling.** Uniformly sample $o(n)$ columns of $\mathbf{K}$ to form $\mathbf{K}''$; then sample c_3 columns of $\mathbf{K}''$ to construct $\mathbf{C}_3$ using the adaptive sampling algorithm (Algorithm 2) according to the residual $\mathbf{K}'' - \mathcal{P}_{[\mathbf{C}_1,\, \mathbf{C}_2]}(\mathbf{K}'')$;
5: **return** $\mathbf{C} = [\mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3]$.

5.4 Discussions

The two algorithms discussed in the above have strong theoretical guarantees, but they are not efficient enough for large-scale applications. First, their time complexities are quadratic in n. Second, they require the full observation of $\mathbf{K}$. In fact, at present no existing column selection algorithm satisfies the three properties simultaneously:

1. the time and memory costs are $\mathcal{O}(n)$;
2. only $\mathcal{O}(n)$ entries of $\mathbf{K}$ need to be observed;
3. relative-error bound holds in expectation or with high probability.

It is interesting to find such an algorithm, and it remains an open problem.

Nevertheless, uniform+adaptive^2 is a promising column selection algorithm for it may be adapted to satisfy the above three properties. The drawback of uniform+adaptive^2 is that computing the adaptive sampling probability costs quadratic time and requires the full observation of $\mathbf{K}$. There may be remedies for this problem. The adaptive-partial algorithm in (Kumar et al., 2012) satisfies the first two properties, but it lacks theoretical analysis. Another possibility is to first uniformly sample $o(n)$ columns and then down-sample to $\mathcal{O}(k/\epsilon)$ columns by adaptive sampling, which we describe in Algorithm 4. In this way, the first two properties can be satisfied, and it may be theoretically explained under the incoherent matrix assumption. We do not implement such heuristics for their theoretical property is completely unknown and they are beyond the scope of this paper.

Table 3: A summary of the datasets for kernel approximation.

Dataset	MNIST	Letters	PenDigit	Cpusmall	Mushrooms
#Instance	60,000	15,000	10,992	8,192	8,124
#Attribute	780	16	16	12	112
γ ($\eta = 0.5$)	≈ 1.50	0.155	0.045	0.031	0.850
γ ($\eta = 0.9$)	≈ 2.30	0.290	0.073	0.057	1.140

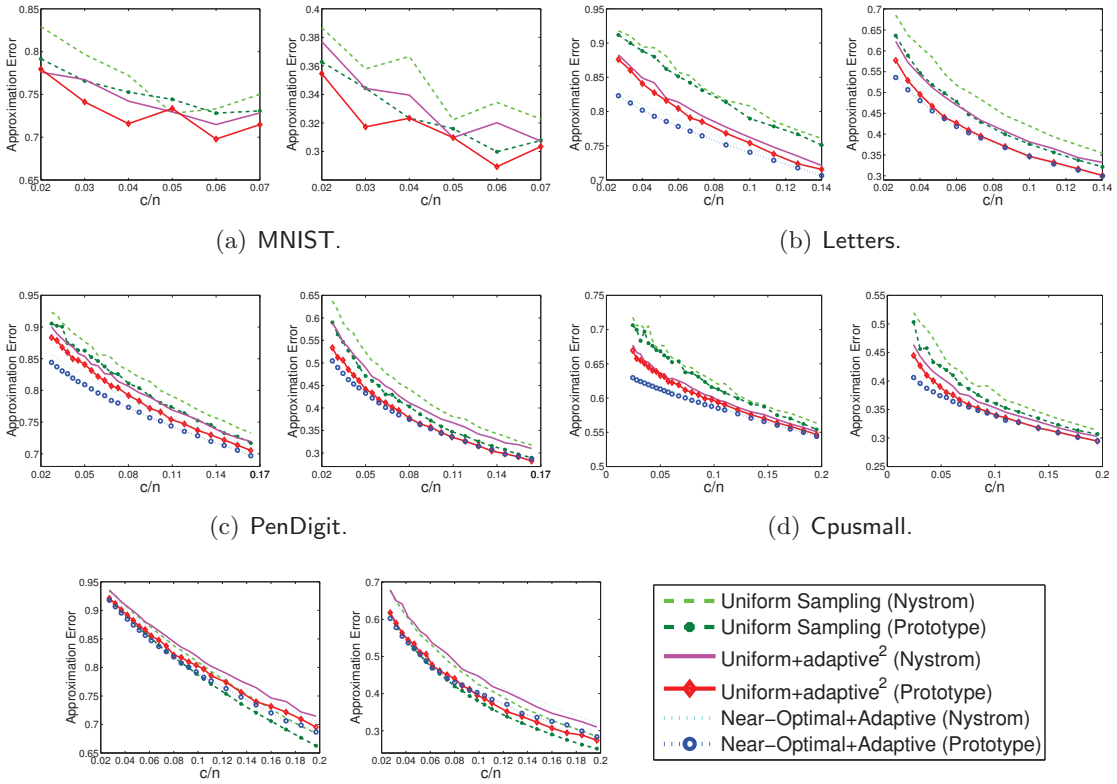

(a) MNIST. (b) Letters.

(c) PenDigit. (d) Cpusmall.

(e) Mushroom. (f) Legend.

Figure 1: The ratio $\frac{c}{n}$ against the error ratio defined in (8). In each subfigure, the left corresponds to the RBF kernel matrix with $\eta = 0.5$, and the right corresponds to $\eta = 0.9$, where η is defined in (7).

6. Experiments on the Column Sampling Algorithms

We empirically conduct comparison among three column selection algorithms—uniform sampling, uniform + adaptive^2, and the near-optimal + adaptive sampling algorithm.

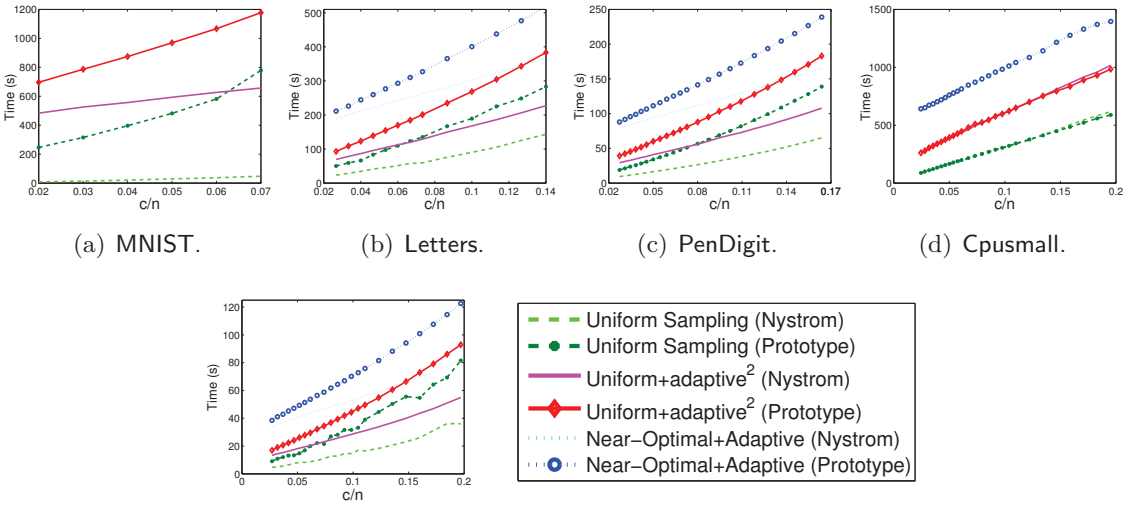

(a) MNIST. (b) Letters. (c) PenDigit. (d) Cpusmall.

(e) Mushroom. (f) Legend.

Figure 2: The growth of the average elapsed time in $\frac{c}{n}$.

6.1 Experiment Setting

We perform experiments on several datasets collected on the LIBSVM website[2] where the data are scaled to [0,1]. We summarize the datasets in Table 3.

For each dataset, we generate a radial basis function (RBF) kernel matrix $\mathbf{K}$ defined by $k_{ij} = \exp(-\frac{1}{2\gamma^2}\|\mathbf{x}_i - \mathbf{x}_j\|_2^2)$. Here $\gamma > 0$ is the scaling parameter; the larger the scaling parameter γ is, the faster the spectrum of the kernel matrix decays (Gittens and Mahoney, 2013). The previous work has shown that for the same dataset, with different settings of γ, the sampling algorithms have very different performances. Instead of setting γ arbitrarily, we set γ in the following way.

Letting $p = \lceil 0.05n \rceil$, we define

$$\eta \triangleq \frac{\sum_{i=1}^{p} \lambda_i^2(\mathbf{K})}{\sum_{i=1}^{n} \lambda_i^2(\mathbf{K})} = \frac{\|\mathbf{K}_p\|_F^2}{\|\mathbf{K}\|_F^2}, \tag{7}$$

which denotes the ratio of the top 5% eigenvalues of the kernel matrix $\mathbf{K}$ to the all eigenvalues. In general, a large γ results in a large η. For each dataset, we use two different settings of γ such that $\eta = 0.5$ or $\eta = 0.9$.

The models and algorithms are all implemented in MATLAB. We run the algorithms on a workstation with Intel Xeon 2.40GHz CPUs, 24GB memory, and 64bit Windows Server 2008 system. To compare the running time, we set MATLAB in single thread mode by the command "maxNumCompThreads(1)". In the experiments we do not keep $\mathbf{K}$ in memory. We use a variant of Algorithm 1—we compute and store one block, instead of one column, of $\mathbf{K}$ at a time. We keep at most $1,000$ columns of $\mathbf{K}$ in memory at a time.

2. http://www.csie.ntu.edu.tw/∼cjlin/libsvmtools/datasets/

We set the target rank k to be $k = \lceil n/100 \rceil$ in all the experiments unless otherwise specified. We evaluate the performance by

$$\text{Approximation Error} \;=\; \|\mathbf{K} - \tilde{\mathbf{K}}\|_F / \|\mathbf{K}\|_F, \tag{8}$$

where $\tilde{\mathbf{K}}$ is the approximation generated by each method.

To evaluate the quality of the approximate rank-k eigenvalue decomposition, we use *misalignment* to indicate the distance between the true eigenvectors $\mathbf{U}_k$ $(n \times k)$ and the approximate eigenvectors $\tilde{\mathbf{V}}_k$ $(n \times k)$:

$$\text{Misalignment} \;=\; \frac{1}{k}\left\|\mathbf{U}_k - \tilde{\mathbf{V}}_k \tilde{\mathbf{V}}_k^T \mathbf{U}_k\right\|_F^2 \;\in\; [0, 1]. \tag{9}$$

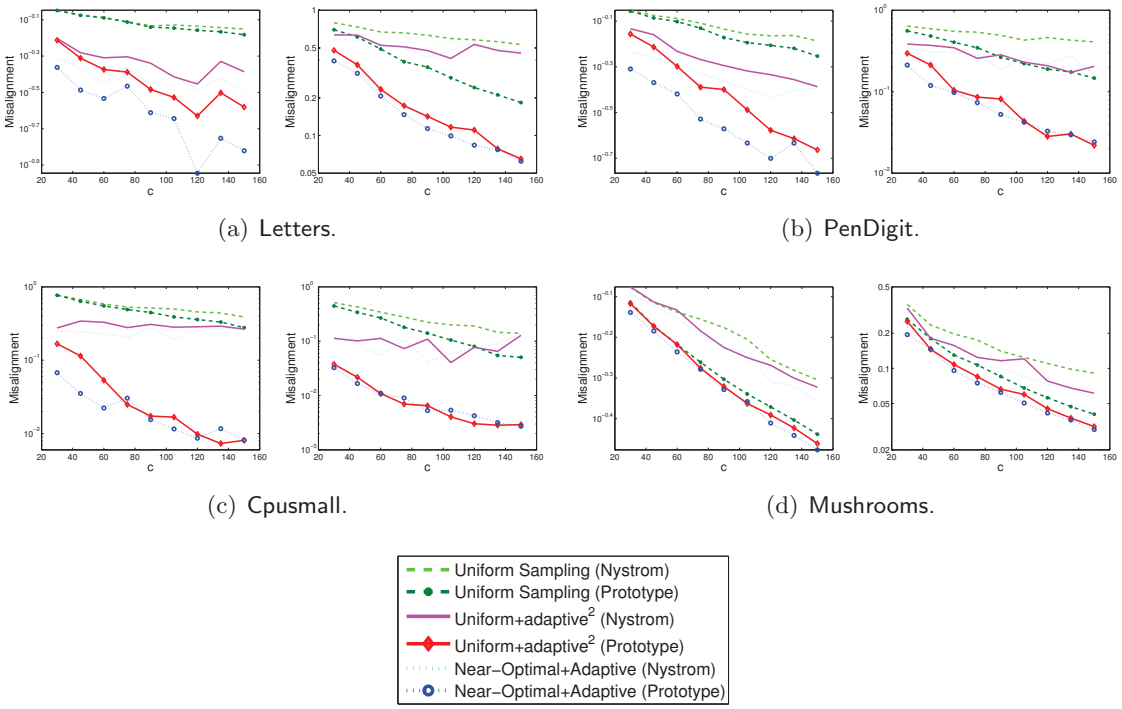

(a) Letters. (b) PenDigit.

(c) Cpusmall. (d) Mushrooms.

(e) Legend.

Figure 3: The number of selected columns c against the misalignment (log-scale) defined in (9). In each subfigure, the left corresponds to the RBF kernel matrix with $\eta = 0.5$, and the right corresponds to $\eta = 0.9$, where η is defined in (7).

6.2 Matrix Approximation Accuracy

In the first set of experiments, we compare the matrix approximation quality using the Frobenius norm approximation error defined in (8) as the metric.

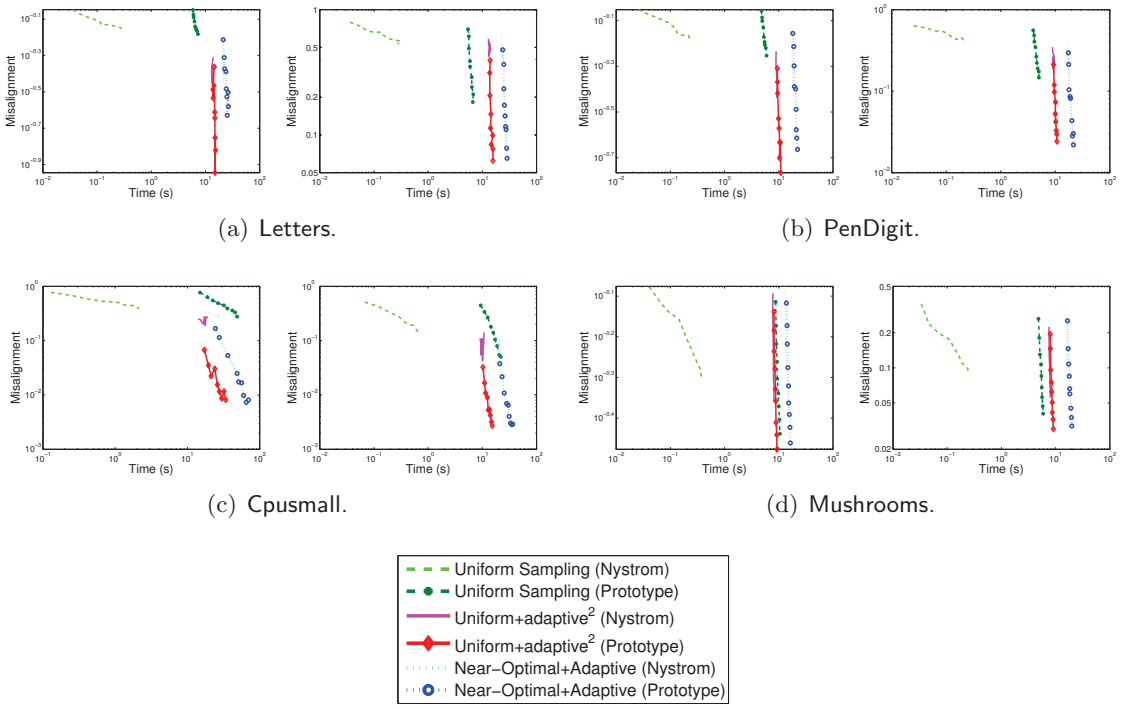

(a) Letters. (b) PenDigit.

(c) Cpusmall. (d) Mushrooms.

(e) Legend.

Figure 4: The elapsed time (log-scale) against the misalignment (log-scale) defined in (9). In each subfigure, the left corresponds to the RBF kernel matrix with $\eta = 0.5$, and the right corresponds to $\eta = 0.9$, where η is defined in (7).

Every time when we do column sampling, we repeat each sampling algorithm 10 times and record the minimal approximation error of the 10 repeats. We report the average elapsed time of the 10 repeat rather than the total elapsed time because the 10 repeats can be done in parallel on 10 machines. We plot c against the approximation error in Figure 1. For the kernel matrices with $\eta = 0.9$, we plot c against the average elapsed time in Figure 2; for $\eta = 0.5$, the curve of the running time is very similar, so we do not show it.

The results show that our uniform+adaptive^2 algorithm achieves accuracy comparable with the near-optimal+adaptive algorithm. Especially, when c is large, these two algorithms have virtually the same accuracy, which agrees with our analysis: a large c implies a small error term ϵ, and the error bounds of the two algorithms coincide when ϵ is small. As for the running time, we can see that our uniform+adaptive^2 algorithm is much more efficient than the near-optimal+adaptive algorithm.

Particularly, the MNIST dataset has $60,000$ instances, and the $60,000 \times 60,000$ kernel matrix $\mathbf{K}$ does not fit in memory. The experiment shows that neither the prototype model nor the uniform+adaptive^2 algorithm require keeping $\mathbf{K}$ in memory.

6.3 Kernel Principal Component Analysis

In the second set of experiment, we apply the kernel approximation methods to approximately compute the rank $k = 3$ eigenvalue decomposition of the RBF kernel matrix. We use the *misalignment* defined in (9) as the metric, which reflects the distance between the true and the approximate eigenvectors. We report the average misalignment of the 20 repeats. We do not conduct experiments on the MNIST dataset because the true eigenvectors are too expensive to compute.

To evaluate the memory efficiency, we plot c against the misalignment in Figure 3. The results show that the two non-uniform sampling algorithms are significantly better than uniform sampling. The performance of our uniform+adaptive^2 algorithm is nearly the same to the near-optimal+adaptive algorithm.

To evaluate the time efficiency, we plot the elapsed time against the misalignment in Figure 4. Though the uniform sampling algorithm is the most efficient in most cases, its accuracy is unsatisfactory. In terms of time efficiency, the uniform+adaptive^2 algorithm is better than the near-optimal+adaptive algorithm.

The experiment on the kernel PCA shows that the prototype model with the uniform + adaptive^2 column sampling algorithm achieves the best performance. Though the Nyström method with uniform sampling is the most efficient, its resulting misalignment is worse by an order of magnitude. Therefore, when applied to speedup eigenvalue decomposition, the Nyström method may not be a good choice, especially when high accuracy is required.

6.4 Separating the Time Costs

Besides the total elapsed time, the readers may be interested in the time cost of each step, especially when the data do no fit in memory. We run the Nyström method and the prototype model, each with uniform+adaptive^2 column sampling algorithm, on the MNIST dataset with $\gamma = 2.3$ (see Table 3). Notice that the $60,000 \times 60,000$ kernel matrix does not fit in memory, so we keep at most $1,000$ columns of the kernel matrix in memory at a time. We set $c = 50$ or 500 and repeat the procedure 20 times and record the average elapsed time. In Figure 5 we separately show the time costs of the uniform+adaptive^2 algorithm and the computation of the intersection matrices. In addition, we separate the time costs of evaluating the kernel functions and all the other computations (e.g. SVD of $\mathbf{C}$ and matrix multiplications).

We can see from Figure 5 that when $\mathbf{K}$ does not fit in memory, the computation of the kernel matrix contributes to the most of the computations. By comparing the two subfigures in Figure 5, we can see that as c increases, the costs of computing the kernel matrix barely change, but the costs of other matrix operations significantly increase.

7. The Spectral Shifting Model

All the low-rank approximation methods work well only when the bottom eigenvalues of $\mathbf{K}$ are near zero. In this section we develop extensions of the three SPSD matrix approximation models to tackle matrices with relatively big bottom eigenvalues. We call the proposed method the *spectral shifting (SS) model* and describe it in Algorithm 5. We show that the SS model has stronger error bound than the prototype model.

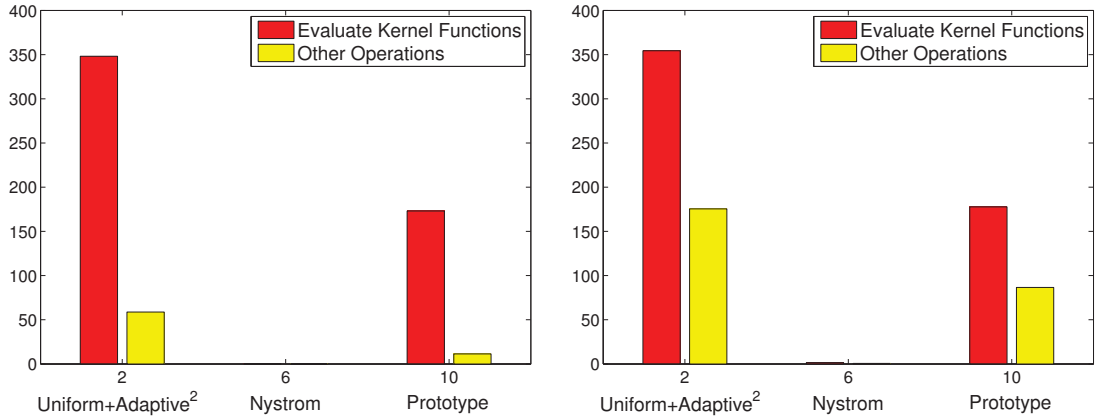

(a) $c = 50$ (b) $c = 500$

Figure 5: The time costs (s) of the uniform+adaptive^2 algorithm and the computation of the intersection matrices.

Algorithm 5 The Spectral Shifting Method.

Input: an $n \times n$ SPSD matrix $\mathbf{K}$, a target rank k, the number of sampled columns c, the oversampling parameter l.
// (optional) approximately do the initial spectral shifting
$\mathbf{\Omega} \longleftarrow n \times l$ standard Gaussian matrix;
$\mathbf{Q} \longleftarrow$ the l orthonormal basis of $\mathbf{K\Omega} \in \mathbb{R}^{n\times l}$;
$s \longleftarrow$ sum of the top k singular values of $\mathbf{Q}^T\mathbf{K} \in \mathbb{R}^{l\times n}$;
$\tilde{\delta} = \frac{1}{n-k}\big(\mathrm{tr}(\mathbf{K}) - s\big) \approx \bar{\delta}$;
$\bar{\mathbf{K}} \leftarrow \mathbf{K} - \tilde{\delta}\mathbf{I}_n \in \mathbb{R}^{n\times n}$;
// perform sketching, e.g. random projection or column selection
$\bar{\mathbf{C}} = \bar{\mathbf{K}}\mathbf{P}$, where $\mathbf{P}$ is an $n \times c$ random projection or selection matrix;
Optional: replace $\bar{\mathbf{C}}$ by its orthonormal bases;
// compute the spectral shifting parameter and the intersection matrix
$\delta^{\mathrm{ss}} \longleftarrow \frac{1}{n-\mathrm{rank}(\bar{\mathbf{C}})}\Big(\mathrm{tr}(\mathbf{K}) - \mathrm{tr}\big(\bar{\mathbf{C}}^\dagger\mathbf{K}\bar{\mathbf{C}}\big)\Big)$;
$\mathbf{U}^{\mathrm{ss}} \longleftarrow \bar{\mathbf{C}}^\dagger\mathbf{K}(\bar{\mathbf{C}}^\dagger)^T - \delta^{\mathrm{ss}}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger$;
return the approximation $\tilde{\mathbf{K}}_c^{\mathrm{ss}} = \bar{\mathbf{C}}\mathbf{U}^{\mathrm{ss}}\bar{\mathbf{C}}^T + \delta^{\mathrm{ss}}\mathbf{I}_n$.

In Section 7.1 we formulate the SS model. In Section 7.2 we study SS from an optimization perspective. In Section 7.3 we show that SS has better error bound than the prototype model. Especially, with the near-optimal+adaptive column sampling algorithm, SS demonstrates much stronger error bound than the existing matrix approximation methods. In Section 7.4 we provide an efficient algorithm for computing the initial spectral shifting term. In Section 7.5 we discuss how to combine spectral shifting with other kernel approximation methods.

7.1 Model Formulation

The spectral shifting (SS) model is defined by

$$\tilde{\mathbf{K}}_c^{\text{ss}} \;=\; \bar{\mathbf{C}}\mathbf{U}^{\text{ss}}\bar{\mathbf{C}}^T + \delta^{\text{ss}}\mathbf{I}_n. \tag{10}$$

Here $\delta^{\text{ss}} \geq 0$ is called the spectral shifting term. This approximation is computed in three steps. Firstly, (approximately) compute the initial spectral shifting term

$$\bar{\delta} \;=\; \frac{1}{n-k}\Big(\text{tr}(\mathbf{K}) - \sum_{j=1}^{k}\sigma_j(\mathbf{K})\Big), \tag{11}$$

and then perform spectral shifting $\bar{\mathbf{K}} = \mathbf{K} - \bar{\delta}\mathbf{I}_n$, where $k \leq c$ is the target rank. This step is optional. Due to Theorem 6 and Remark 7, SS is better than the prototype model even if $\bar{\delta} = 0$; however, without this step, the quantity of the improvement contributed by SS is unknown. Secondly, draw a column selection matrix $\mathbf{P}$ and form the sketch $\bar{\mathbf{C}} = \bar{\mathbf{K}}\mathbf{P}$. Finally, with $\bar{\mathbf{C}}$ at hand, compute δ^{ss} and $\mathbf{U}^{\text{ss}}$ by

$$\begin{aligned} \delta^{\text{ss}} \;&=\; \frac{1}{n - \text{rank}(\bar{\mathbf{C}})}\Big(\text{tr}(\mathbf{K}) - \text{tr}\big(\bar{\mathbf{C}}^\dagger\mathbf{K}\bar{\mathbf{C}}\big)\Big), \\ \mathbf{U}^{\text{ss}} \;&=\; \bar{\mathbf{C}}^\dagger\mathbf{K}(\bar{\mathbf{C}}^\dagger)^T - \delta^{\text{ss}}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger. \end{aligned} \tag{12}$$

We will show that $\tilde{\mathbf{K}}_c^{\text{ss}}$ is positive (semi)definite if $\mathbf{K}$ is positive (semi)definite.

However, when the bottom eigenvalues of $\mathbf{K}$ are small, the computed spectral shifting term is small, where there is little difference between SS and the prototype model, and the spectral shifting operation is not advised.

7.2 Optimization Perspective

The SS model is an extension of the prototype model from the optimization perspective. Given an SPSD matrix $\mathbf{K}$, the prototype model computes the sketch $\mathbf{C} = \mathbf{K}\mathbf{P}$ and the intersection matrix

$$\mathbf{U}^\star \;=\; \mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T \;=\; \underset{\mathbf{U}}{\text{argmin}}\,\|\mathbf{K} - \mathbf{C}\mathbf{U}\mathbf{C}^T\|_F^2. \tag{13}$$

Analogously, with the sketch $\bar{\mathbf{C}} = \bar{\mathbf{K}}\mathbf{P} = \mathbf{K}\mathbf{P} - \bar{\delta}\mathbf{P}$ at hand, SS is obtained by solving

$$\big(\mathbf{U}^{\text{ss}}, \delta^{\text{ss}}\big) \;=\; \underset{\mathbf{U},\delta}{\text{argmin}}\,\big\|\mathbf{K} - \bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T - \delta\mathbf{I}_n\big\|_F^2, \tag{14}$$

obtaining the intersection matrix $\mathbf{U}^{\text{ss}}$ and the spectral shifting term δ^{ss}. By analyzing the optimization problem (14), we obtain the following theorem. Its proof is in Appendix E.

Theorem 6 *The pair $(\delta^{ss}, \mathbf{U}^{ss})$ defined in (12) is the global minimizer of problem (14), which indicates that using any other $(\delta, \mathbf{U})$ to replace $(\delta^{ss}, \mathbf{U}^{ss})$ results in a larger approximation error. Furthermore, if $\mathbf{K}$ is positive (semi)definite, then the approximation $\bar{\mathbf{C}}\mathbf{U}^{ss}\bar{\mathbf{C}}^T + \delta^{ss}\mathbf{I}_n$ is also positive (semi)definite.*

Remark 7 *The optimization perspective indicates the superiority of SS. Suppose we skip the initial spectral shifting step and simply set $\bar{\delta} = 0$. Then $\bar{\mathbf{C}} = \mathbf{C}$. If the constraint $\delta = 0$ is to the optimization problem (14), then (14) will become identical to the prototype model (13). Obviously, adding this constraint will make the optimal objective function value get worse, so the optimal objective function value of (14) is always less than or equal to (13). Hence, without the initial spectral shifting step, SS is still more accurate than the prototype model.*

7.3 Error Analysis

The following theorem indicates that the SS model with any spectral shifting term $\delta \in (0, \bar{\delta}]$ has a stronger bound than the prototype model. The proof is in Appendix F.

Theorem 8 *Suppose there is a sketching matrix $\mathbf{P} \in \mathbb{R}^{n\times c}$ such that for any $n \times n$ symmetric matrix $\mathbf{A}$ and target rank k ($\ll n$), by forming $\mathbf{C} = \mathbf{AP}$, the prototype model satisfies the error bound*

$$\left\|\mathbf{A} - \mathbf{CC}^{\dagger}\mathbf{A}(\mathbf{C}^{\dagger})^{T}\mathbf{C}^{T}\right\|_F^2 \;\leq\; \eta \left\|\mathbf{A} - \mathbf{A}_k\right\|_F^2$$

for certain $\eta > 0$. Let $\mathbf{K}$ be any $n \times n$ SPSD matrix, $\tilde{\delta} \in (0, \bar{\delta}]$ be the initial spectral shifting term where $\bar{\delta}$ is defined in (11), $\bar{\mathbf{K}} = \mathbf{K} - \tilde{\delta}\mathbf{I}_n$, $\bar{\mathbf{C}} = \bar{\mathbf{K}}\mathbf{P}$, and $\tilde{\mathbf{K}}_c^{ss}$ be the SS model defined in (10). Then

$$\left\|\mathbf{K} - \tilde{\mathbf{K}}_c^{ss}\right\|_F^2 \;\leq\; \eta \left\|\bar{\mathbf{K}} - \bar{\mathbf{K}}_k\right\|_F^2 \;\leq\; \eta \left\|\mathbf{K} - \mathbf{K}_k\right\|_F^2.$$

We give an example in Figure 6 to illustrate the intuition of spectral shifting. We use the toy matrix $\mathbf{K}$: an $n\times n$ SPSD matrix whose the t-th eigenvalue is 1.05^{-t}. We set $n = 100$ and $k = 30$, and hence $\bar{\delta} = 0.064$. From the plot of the eigenvalues we can see that the "tail" of the eigenvalues becomes thinner after the spectral shifting. Specifically, $\|\mathbf{K} - \mathbf{K}_k\|_F^2 = 0.52$ and $\|\bar{\mathbf{K}} - \bar{\mathbf{K}}_k\|_F^2 \leq 0.24$. From Theorem 8 we can see that if $\|\mathbf{K} - \mathbf{CC}^{\dagger}\mathbf{K}(\mathbf{C}^{\dagger})^{T}\mathbf{C}^{T}\|_F \leq 0.52\eta$, then $\|\mathbf{K} - \tilde{\mathbf{K}}_c^{\mathrm{ss}}\|_F \leq 0.24\eta$. This indicates that SS has much stronger error bound than the prototype model.

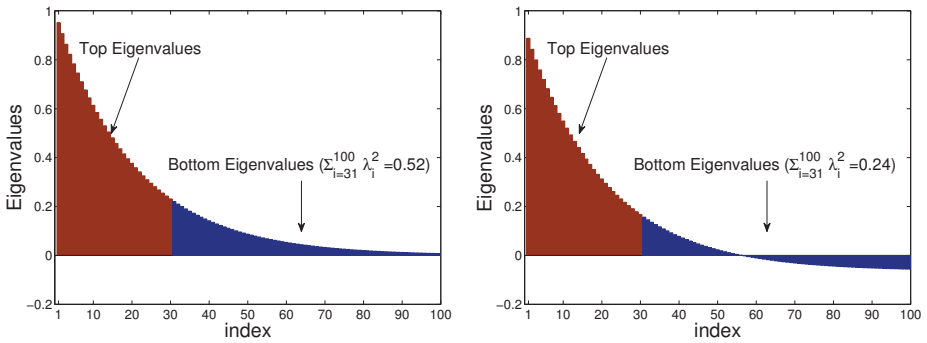

(a) Before spectral shifting. (b) After spectral shifting.

Figure 6: We plot the eigenvalues of $\mathbf{K}$ in Figure 6(a) and $\bar{\mathbf{K}} = \mathbf{K} - \bar{\delta}\mathbf{I}_{100}$ in Figure 6(b).

The following theorem shows an error bound of the SS model, which is stronger than the prototype model, especially when the bottom $n-k$ eigenvalues of $\mathbf{K}$ are big. The proof is in Appendix F.

Theorem 9 *Suppose there is a sketching matrix $\mathbf{P} \in \mathbb{R}^{n\times c}$ such that for any $n \times n$ symmetric matrix $\mathbf{A}$ and target rank k ($\ll n$), by forming the sketch $\mathbf{C} = \mathbf{AP}$, the prototype model satisfies the error bound*

$$\left\|\mathbf{A} - \mathbf{C}\mathbf{C}^{\dagger}\mathbf{A}(\mathbf{C}^{\dagger})^T\mathbf{C}^T\right\|_F^2 \;\leq\; \eta\left\|\mathbf{A}-\mathbf{A}_k\right\|_F^2$$

for certain $\eta > 0$. Let $\mathbf{K}$ be any $n \times n$ SPSD matrix, $\bar{\delta}$ defined in (11) be the initial spectral shifting term, and $\tilde{\mathbf{K}}_c^{ss}$ be the SS model defined in (10). Then

$$\left\|\mathbf{K} - \tilde{\mathbf{K}}_c^{ss}\right\|_F^2 \;\leq\; \eta\left(\left\|\mathbf{K}-\mathbf{K}_k\right\|_F^2 - \frac{\left[\sum_{i=k+1}^{n}\lambda_i(\mathbf{K})\right]^2}{n-k}\right).$$

If $\bar{\mathbf{C}}$ contains the columns of $\bar{\mathbf{K}}$ sampled by the near-optimal+adaptive algorithm in Theorem 3, which has the strongest bound, then the error bound incurred by SS is given in the following corollary.

Corollary 10 *Suppose we are given any SPSD matrix $\mathbf{K}$ and we sample $c = \mathcal{O}(k/\epsilon)$ columns of $\bar{\mathbf{K}}$ to form $\bar{\mathbf{C}}$ using the near-optimal+adaptive column sampling algorithm (Theorem 3). Then the inequality holds:*

$$\mathbb{E}\left\|\mathbf{K} - \tilde{\mathbf{K}}_c^{ss}\right\|_F^2 \leq (1+\epsilon)\left(\left\|\mathbf{K}-\mathbf{K}_k\right\|_F^2 - \frac{\left[\sum_{i=k+1}^{n}\lambda_i(\mathbf{K})\right]^2}{n-k}\right).$$

Here we give an example to demonstrate the superiority of SS over the prototype model, the Nyström method, and even the truncated SVD of the same scale.

Example 1 *Let $\mathbf{K}$ be an $n \times n$ SPSD matrix such that $\lambda_1(\mathbf{K}) \geq \cdots \geq \lambda_k(\mathbf{K}) > \theta = \lambda_{k+1}(\mathbf{K}) = \cdots = \lambda_n(\mathbf{K}) > 0$. By sampling $c = \mathcal{O}(k)$ columns by the near-optimal+adaptive algorithm (Theorem 3), we have that*

$$\left\|\mathbf{K} - \tilde{\mathbf{K}}_c^{ss}\right\|_F^2 \;=\; 0$$

and that

$$(n-c)\theta^2 \;=\; \left\|\mathbf{K}-\mathbf{K}_c\right\|_F^2 \;\leq\; \left\|\mathbf{K}-\tilde{\mathbf{K}}_c^{proto}\right\|_F^2 \;\leq\; \left\|\mathbf{K}-\tilde{\mathbf{K}}_c^{nys}\right\|_F^2.$$

Here $\tilde{\mathbf{K}}_c^{proto}$ and $\tilde{\mathbf{K}}_c^{nys}$ respectively denote the approximation formed by the prototype model and the Nyström method. In this example the SS model is far better than the other models if we set θ as a large constant.

7.4 Approximately Computing $\bar{\delta}$

The SS model uses $\bar{\delta}$ as the initial spectral shifting term. However, computing $\bar{\delta}$ according to (11) requires the partial eigenvalue decomposition which costs $\mathcal{O}(n^2k)$ time and $\mathcal{O}(n^2)$ memory. For large-scale data, one can simply set $\bar{\delta} = 0$; Remark 7 shows that SS with this setting still works better than the prototype model. For medium-scale data, one can approximately compute $\bar{\delta}$ by the algorithm devised and analyzed in this subsection.

We depict the algorithm in Lines 2–6 of Algorithm 5. The performance of the approximation is analyzed in the following theorem.

Theorem 11 *Let $\bar{\delta}$ be defined in (11) and $\tilde{\delta}$, k, l, n be defined in Algorithm 5. The following inequality holds:*

$$\mathbb{E}\left[\left|\bar{\delta} - \tilde{\delta}\right| / \bar{\delta}\right] \;\leq\; k/\sqrt{l},$$

where the expectation is taken w.r.t. the Gaussian random matrix $\mathbf{\Omega}$ in Algorithm 5. Lines 2–6 in Algorithm 5 compute $\tilde{\delta}$ in $\mathcal{O}(n^2l)$ time and $\mathcal{O}(nl)$ memory.

Here we empirically evaluate the accuracy of the approximation to $\bar{\delta}$ (Lines 2–6 in Algorithm 5) proposed in Theorem 11. We use the RBF kernel matrices with the scaling parameter γ listed Table 3. We use the error ratio $|\bar{\delta} - \tilde{\delta}|/\bar{\delta}$ to evaluate the approximation quality. We repeat the experiments 20 times and plot l/k against the average error ratio in Figure 7. Here $\tilde{\delta}$, l, and k are defined in Theorem 11. We can see that the approximation of $\bar{\delta}$ has high quality: when $l = 4k$, the error ratios are less than 0.03 in all cases, no matter whether the spectrum of $\mathbf{K}$ decays fast or slow.

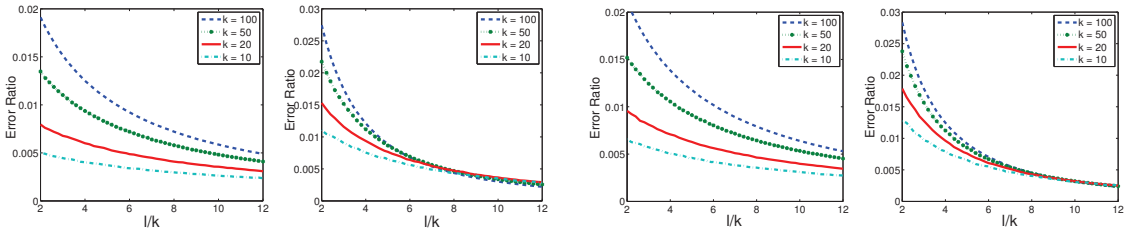

(a) Letters (b) PenDigit

Figure 7: The ratio $\frac{l}{k}$ against the error $|\bar{\delta} - \tilde{\delta}|/\bar{\delta}$. In each subfigure, the left corresponds to the RBF kernel matrix with $\eta = 0.5$, and the right corresponds to $\eta = 0.9$, where η is defined in (7).

7.5 Combining with Other Matrix Approximation Methods

There are many other matrix approximation approaches such as the ensemble Nyström method (Kumar et al., 2012) and MEKA (Si et al., 2014). In fact, the key components of the ensemble Nyström method and MEKA are the Nyström method, which can be straightforwardly replaced by other matrix approximation methods such as the SS model.

The ensemble Nyström method improves the Nyström method by running the Nyström method t times and combine the samples to construct the kernel approximation:

$$\tilde{\mathbf{K}}^{\text{ens}}_{t,c} = \sum_{i=1}^{t} \mu^{(i)} \mathbf{C}^{(i)} {\mathbf{W}^{(i)}}^{\dagger} {\mathbf{C}^{(i)}}^{T},$$

where $\mu^{(1)}, \cdots, \mu^{(t)}$ are the weights of the samples, and a simple but effective strategy is to set the weights as $\mu^{(1)} = \cdots = \mu^{(t)} = \frac{1}{t}$. However, the time and memory costs of computing $\mathbf{C}$ and $\mathbf{U}$ are respectively t times as much as that of its base method (e.g. the Nyström method), and the ensemble Nyström method needs to use the Sherman-Morrison-Woodbury formula t times to combine the samples. When executed on a single machine, the accuracy gained by the ensemble may not worth the t times more time and memory costs.

MEKA is reported to be the state-of-the-art kernel approximation method. It exploits the block-diagonal structure of kernel matrices, and outputs an $n \times c$ sparse matrix $\mathbf{C}$ and a $c \times c$ small matrix $\mathbf{U}$ such that $\mathbf{K} \approx \mathbf{C}\mathbf{U}\mathbf{C}^T$. MEKA first finds the blocks by clustering the data into b clusters and permutes the kernel matrix accordingly. It then approximates the diagonal blocks by the Nyström method, *which can be replaced by other methods.* It finally approximates the off-diagonal blocks using the diagonal blocks. If the kernel matrix is partitioned into b^2 blocks, only b blocks among the b^2 blocks of $\mathbf{C}$ are nonzero, and the number of nonzero entries of $\mathbf{C}$ is at most $\text{nnz}(\mathbf{C}) = nc/b$. MEKA is thus much more memory efficient than the Nyström method. If we use the SS model to approximate the diagonal blocks, then the resulting MEKA approximation will be in the form $\mathbf{K} \approx \mathbf{C}\mathbf{U}\mathbf{C}^T + \delta_1 \mathbf{I} \oplus \cdots \oplus \delta_b \mathbf{I}$, where δ_i corresponds to the i-th diagonal block.

8. Empirically Evaluating the Spectral Shifting Model

We empirically evaluate the spectral shifting method by comparing the following kernel approximation models in terms of approximation quality and the generalization performance on Gaussian process regression.

- The Nyström method with the uniform+adaptive^2 algorithm.
- The prototype model with the uniform+adaptive^2 algorithm.
- The memory efficient kernel approximation (MEKA) method (Si et al., 2014), which approximates the diagonal blocks of the kernel matrix by the Nyström method. We use the code released by the authors with default settings.
- The spectral shifting (SS) model with the uniform+adaptive^2 algorithm.
- SS+MEKA: the same to MEKA except for using SS, rather than the Nyström method, to approximate the diagonal blocks.

Since the experiments are all done on a single machine, the time and memory costs of the ensemble Nyström method (Kumar et al., 2012) are t times larger. It would be unfair to directly do comparison with the ensemble Nyström method.

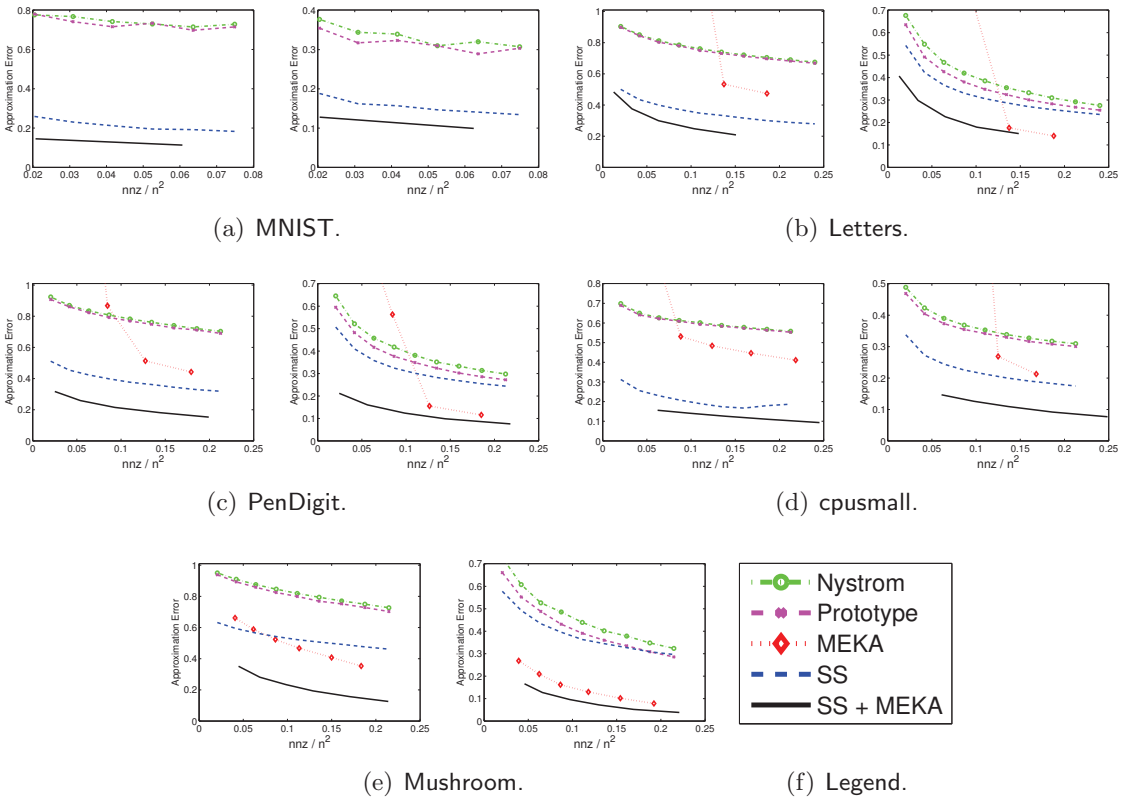

(a) MNIST. (b) Letters.

(c) PenDigit. (d) cpusmall.

(e) Mushroom. (f) Legend.

Figure 8: The memory cost against approximation error. Here "nnz" is number of nonzero entries in the sketch, namely, $\text{nnz}(\mathbf{C}) + \text{nnz}(\mathbf{U})$. In each subfigure, the left corresponds to the RBF kernel matrix with $\eta = 0.5$, and the right corresponds to $\eta = 0.9$, where η is defined in (7).

8.1 Experiments on Approximation Accuracy

We conduct experiments using the same setting as in Section 6.1. To demonstrate the effect of spectral shifting, we only consider a kernel matrix with slowly decaying spectrum. Thus we set η (defined in (7)) relatively small: $\eta = 0.5$ and $\eta = 0.9$. When η is near one, the spectral shifting term becomes near zero, and spectral shifting makes no difference.

We present the results in two ways: (1) Figure 8 plots the memory usage against the approximation error, where the memory usage is proportional to the number of nonzero entries of $\mathbf{C}$ and $\mathbf{U}$; (2) Figure 9 plots the elapsed time against the approximation error. The results have the following implications.

- Using the spectral shifting technique is better than without using it: SS and SS+MEKA are more accurate than the prototype model and MEKA, respectively. The advantage

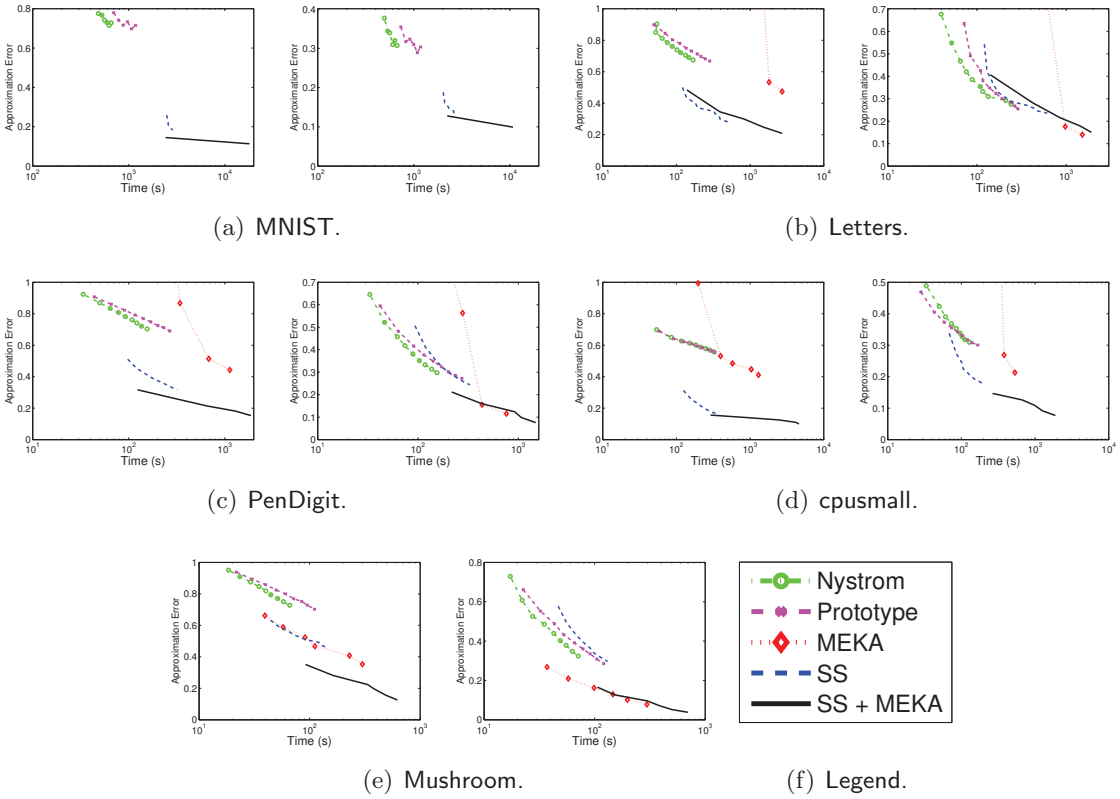

(a) MNIST. (b) Letters.

(c) PenDigit. (d) cpusmall.

(e) Mushroom. (f) Legend.

Figure 9: The elapsed time (log-scale) against the approximation error. In each subfigure, the left corresponds to the RBF kernel matrix with $\eta = 0.5$, and the right corresponds to $\eta = 0.9$, where η is defined in (7).

of spectral shifting is particularly obvious when the spectrum of $\mathbf{K}$ decays slowly, i.e. when $\eta = 0.5$.

- SS and SS+MEKA are the best among the compared methods according to our experiments. Especially, if a high-quality approximation in limited memory is desired, SS+MEKA should be the best choice.

- The kernel matrix of the MNIST dataset is $60,000 \times 60,000$, which does not fit in the 24GB memory. This shows that the compared methods and sampling algorithm do not require keeping $\mathbf{K}$ in memory.

In addition, we conduct experiments on the large-scale dataset—covtype—which has 581,012 data instances. We compare among the Nyström method, MEKA, and MEKA+SS. The experiment setting is slightly different from the ones in the above. For each of the four methods, we use uniform sampling to select columns. As for the MEKA based methods, we set the number of clusters to be 30 (whereas the default is 10) to increase scalability. As

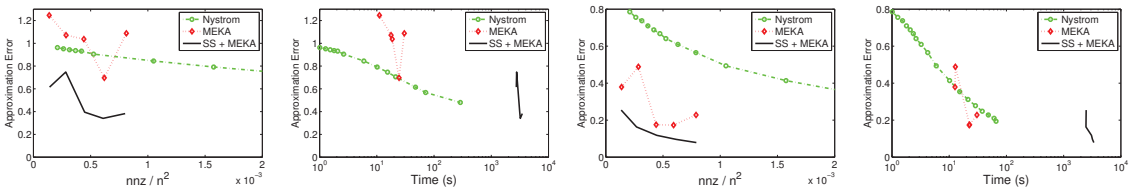

(a) $\eta \approx 0.5$ (b) $\eta \approx 0.9$

Figure 10: The experiments on the covtype dataset. We plot the memory cost against approximation error and plot the elapsed time (log-scale) against the approximation error.

Table 4: Summary of datasets for Gaussian process regression

	Plant	White Wine	Red Wine	Concrete	Energy (Heat)	Energy (Cool)	Housing
#Instance	9,568	4,898	1,599	1,030	768	768	506
#Attribute	4	11	11	8	8	8	13
γ	0.1	1	1	1	0.5	0.5	1
σ^2	0.1	0.01	0.01	0.0002	0.0005	0.001	0.005

for the spectral shifting method, we do not perform the initial spectral shifting. The results are plotted in Figure 10. The results show the effectiveness and scalability of the spectral shift method.

8.2 Experiments on Gaussian Process Regression

We apply the kernel approximation methods to Gaussian process regression (GPR). We assume that the training set is $\{(\mathbf{x}_1, y_1), ..., (\mathbf{x}_n, y_n)\}$, where $\mathbf{x}_i \in \mathbb{R}^d$ are input vectors and $y_i \in \mathbb{R}$ are the corresponding outputs. GPR is defined as

$$y = u + f(\mathbf{x}) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2),$$

where $f(\mathbf{x})$ follows a Gaussian process with mean function 0 and kernel function $\kappa(\cdot, \cdot)$. Furthermore, we define the kernel matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$, where the (i, j)-th entry of $\mathbf{K}$ is $\kappa(\mathbf{x}_i, \mathbf{x}_j)$.

For a test input $\mathbf{x}_*$, the prediction is given by

$$\hat{y}_* = \mathbf{k}^T(\mathbf{x}_*)(\mathbf{K} + \sigma^2\mathbf{I})^{-1}\mathbf{y},$$

where $\mathbf{y} = [y_1, \ldots, y_n]^T$ and $\mathbf{k}(\mathbf{x}_*) = [\kappa(\mathbf{x}_*, \mathbf{x}_1), \ldots, \kappa(\mathbf{x}_*, \mathbf{x}_n)]^T$. We apply different kernel approximation methods to approximate $\mathbf{K}$ and approximately compute $(\mathbf{K} + \sigma^2\mathbf{I})^{-1}\mathbf{y}$ according to Section 3.1. We evaluate the generalization performance using the mean squared error:

$$\text{MSE} = \frac{1}{m}\sum_{i=1}^{m}(y_{i*} - \hat{y}_{i*})^2,$$

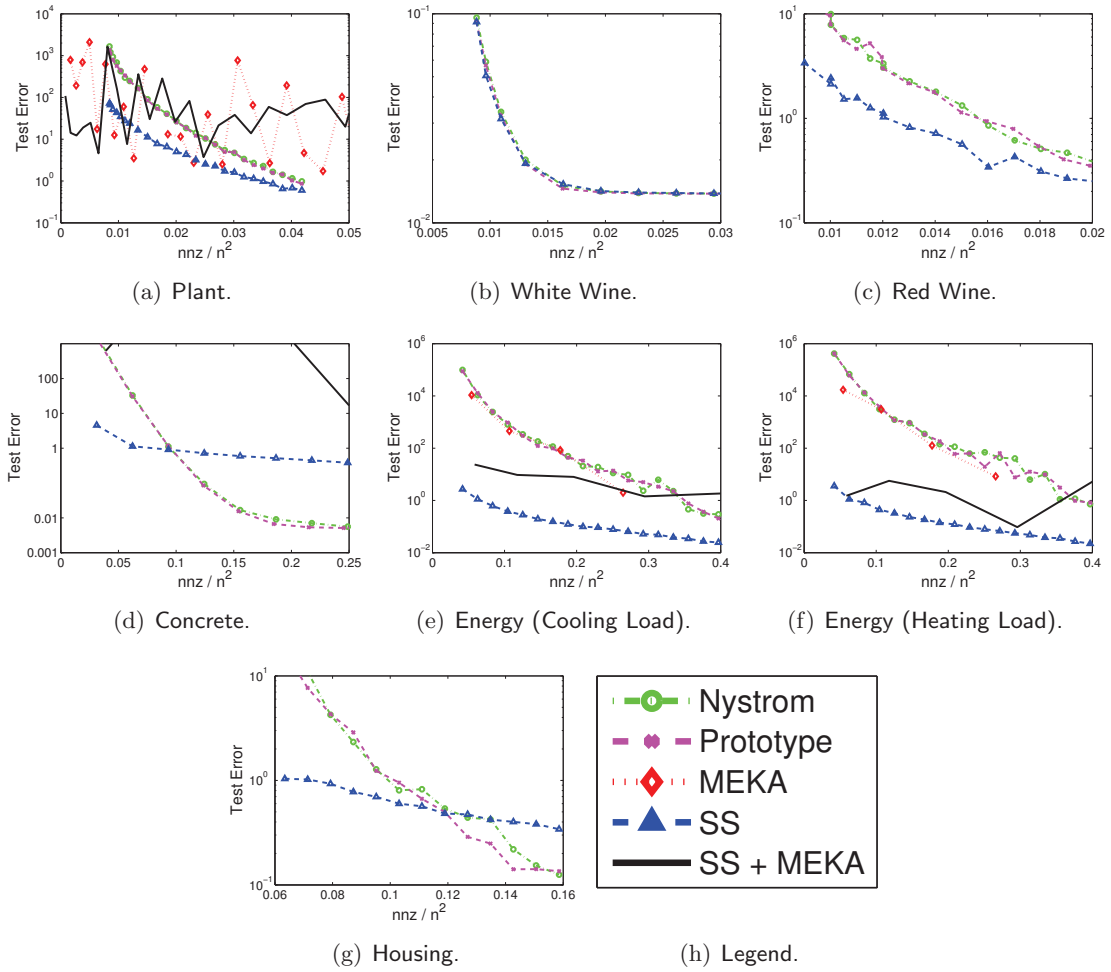

(a) Plant. (b) White Wine. (c) Red Wine.

(d) Concrete. (e) Energy (Cooling Load). (f) Energy (Heating Load).

(g) Housing. (h) Legend.

Figure 11: The results on Gaussian process regression. In the figures "nnz" is number of nonzero entries in the sketch, that is, $\mathrm{nnz}(\mathbf{C}) + \mathrm{nnz}(\mathbf{U})$.

where y_{i*} is the real output of the i-th test sample and m is the number of test samples.

We conduct experiment on seven datasets summarized in Table 4. We use the Gaussian RBF kernel and tune two parameters: the variance σ^2 and the kernel scaling parameter γ. Recall that there are seven compared methods and eight datasets, so the time cost would be daunting if we use cross-validation to find σ and γ for each method on each dataset. Thus we perform a five-fold cross-validation without using kernel approximation to predetermine the two parameters σ and γ, and the same parameters are used for all the kernel approximation methods. We list the obtained parameters in Table 4.

For each of the compared methods, we randomly hold 80% samples for training and the rest for test; we repeat this procedure 50 times and record the average MSE, the average

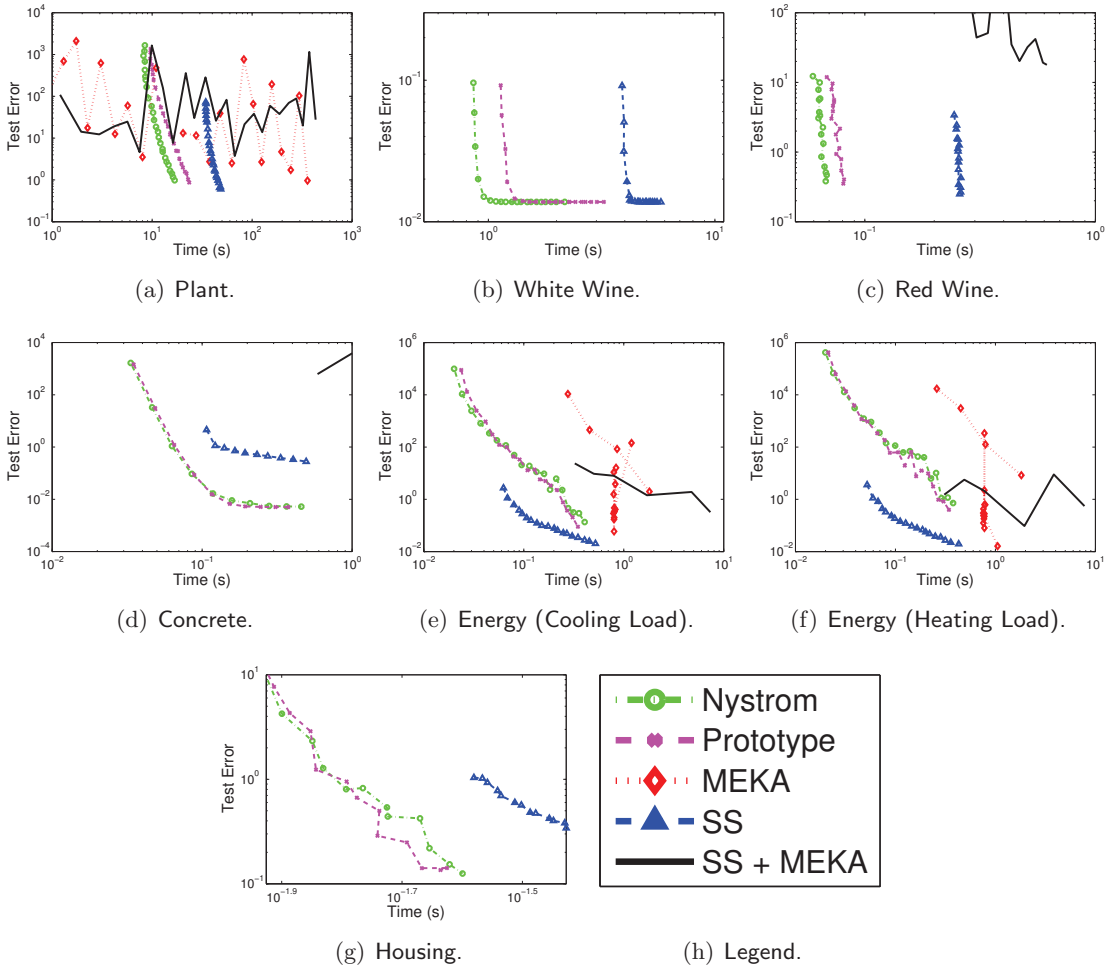

(a) Plant. (b) White Wine. (c) Red Wine.

(d) Concrete. (e) Energy (Cooling Load). (f) Energy (Heating Load).

(g) Housing. (h) Legend.

Figure 12: The results on Gaussian process regression.

elapsed time, and the average of the number of nonzero entries in the sketch. We plot $\frac{\text{nnz}(\mathbf{C})+\text{nnz}(\mathbf{U})}{n^2}$ against MSE in Figure 11 and the elapsed time against MSE in Figure 12.

Using the same amount of memory, our SS model achieves the best performance on the Plant, Red Wine, Energy (Cool), and Energy (Heat) datasets. On the Concrete and Housing datasets, using spectral shifting leads better results unless c is unreasonably large (e.g. $c > 0.1n$). However, using the same amount of time, the compared methods have competitive performance.

MEKA and SS+MEKA in general work very well, but they are numerically unstable on some of the randomly partitioned training data. On the White Wine, Housing, and Concrete datasets, MATLAB occasionally reports errors of numerical instability in approximating the off-diagonal blocks of $\mathbf{K}$ by solving some linear systems; when such happens, we do not report the corresponding test errors in the figures. MEKA and SS+MEKA are also sometimes unstable on the Plant and Red Wine datasets, though MATLAB does not report

error. Although MEKA and SS+MEKA have good performance in general, the numerical instability makes MEKA and SS+MEKA perform very poorly on a few among the 50 randomly partitioned training/test data, and consequently the average test errors become large. This problem can get avoided by repeating MEKA multiple times and choose the one that is stable. However, this will significantly increase the time cost.

9. Conclusions

We have provided an in-depth study of the prototype model for SPSD matrix approximation. First, we have shown that with $c = \mathcal{O}(k/\epsilon)$ columns sampled by the near-optimal+adaptive algorithm, the prototype model attains $1 + \epsilon$ Frobenius norm relative-error bound. This upper bound matches the lower bound up to a constant factor. Second, we have devised a simple column selection algorithm called *uniform+adaptive*^{2}. The algorithm is efficient and very easy to implement, and it has near-optimal relative-error bound. Third, we have proposed an extension called the spectral shifting (SS) model. We have shown that SS has much stronger error bound than the low-rank approximation models, especially when the bottom eigenvalues are not sufficiently small.

Although the prototype model is not very time-efficient, we can resort to the approximate method provided by Wang et al. (2015) to obtain a faster solution to the prototype model. This faster solution requires only linear time and linear memory. The theoretical analysis of the fast solution heavily relies on that of the prototype model, so our established results are useful even if the prototype model itself is not the working horse in real-world applications. In addition, we have shown that the fast solution can also be naturally incorporated with the spectral shifting method.

Acknowledgments

We thank the anonymous reviewers for their helpful suggestions. Wang has been supported by Baidu Scholarship. Luo and Zhang have been supported by the National Natural Science Foundation of China (No. 61572017), Natural Science Foundation of Shanghai City (No. 15ZR1424200), and Microsoft Research Asia Collaborative Research Award.

Appendix A. Proof of Theorem 1

Proof Suppose that $\mathrm{rank}(\mathbf{W}) = \mathrm{rank}(\mathbf{K})$. We have that $\mathrm{rank}(\mathbf{W}) = \mathrm{rank}(\mathbf{C}) = \mathrm{rank}(\mathbf{K})$ because

$$\mathrm{rank}(\mathbf{K}) \;\geq\; \mathrm{rank}(\mathbf{C}) \;\geq\; \mathrm{rank}(\mathbf{W}). \tag{15}$$

Thus there exists a matrix $\mathbf{X}$ such that

$$\begin{bmatrix} \mathbf{K}_{21}^T \\ \mathbf{K}_{22} \end{bmatrix} \;=\; \mathbf{C}\mathbf{X}^T \;=\; \begin{bmatrix} \mathbf{W}\mathbf{X}^T \\ \mathbf{K}_{21}\mathbf{X}^T \end{bmatrix},$$

and it follows that $\mathbf{K}_{21} = \mathbf{X}\mathbf{W}$ and $\mathbf{K}_{22} = \mathbf{K}_{21}\mathbf{X}^T = \mathbf{X}\mathbf{W}\mathbf{X}^T$. Then we have that

$$\mathbf{K} = \begin{bmatrix} \mathbf{W} & (\mathbf{X}\mathbf{W})^T \\ \mathbf{X}\mathbf{W} & \mathbf{X}\mathbf{W}\mathbf{X}^T \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{X} \end{bmatrix} \mathbf{W} \begin{bmatrix} \mathbf{I} & \mathbf{X}^T \end{bmatrix}, \tag{16}$$

$$\mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T = \begin{bmatrix} \mathbf{W} \\ \mathbf{X}\mathbf{W} \end{bmatrix} \mathbf{W}^\dagger \begin{bmatrix} \mathbf{W} & (\mathbf{X}\mathbf{W})^T \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{X} \end{bmatrix} \mathbf{W} \begin{bmatrix} \mathbf{I} & \mathbf{X}^T \end{bmatrix}. \tag{17}$$

Here the second equality in (17) follows from $\mathbf{W}\mathbf{W}^\dagger\mathbf{W} = \mathbf{W}$. We obtain that $\mathbf{K} = \mathbf{C}\mathbf{W}^\dagger\mathbf{C}$. Then we show that $\mathbf{K} = \mathbf{C}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T\mathbf{C}^T$.

Since $\mathbf{C}^\dagger = (\mathbf{C}^T\mathbf{C})^\dagger\mathbf{C}^T$, we have that

$$\mathbf{C}^\dagger = \big(\mathbf{W}(\mathbf{I} + \mathbf{X}^T\mathbf{X})\mathbf{W}\big)^\dagger \mathbf{W}\,[\mathbf{I}, \mathbf{X}^T],$$

and thus

$$\begin{aligned} &\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T\mathbf{W} \\ &= \big(\mathbf{W}(\mathbf{I} + \mathbf{X}^T\mathbf{X})\mathbf{W}\big)^\dagger \mathbf{W}(\mathbf{I} + \mathbf{X}^T\mathbf{X})\Big[\mathbf{W}(\mathbf{I} + \mathbf{X}^T\mathbf{X})\mathbf{W}\big(\mathbf{W}(\mathbf{I} + \mathbf{X}^T\mathbf{X})\mathbf{W}\big)^\dagger\mathbf{W}\Big] \\ &= \big(\mathbf{W}(\mathbf{I} + \mathbf{X}^T\mathbf{X})\mathbf{W}\big)^\dagger \mathbf{W}(\mathbf{I} + \mathbf{X}^T\mathbf{X})\mathbf{W}, \end{aligned}$$

where the second equality follows from Lemma 12 because $(\mathbf{I} + \mathbf{X}^T\mathbf{X})$ is positive definite. Similarly we have

$$\mathbf{W}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T\mathbf{W} = \mathbf{W}\big(\mathbf{W}(\mathbf{I} + \mathbf{X}^T\mathbf{X})\mathbf{W}\big)^\dagger \mathbf{W}(\mathbf{I} + \mathbf{X}^T\mathbf{X})\mathbf{W} = \mathbf{W}.$$

Thus we have

$$\mathbf{C}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T\mathbf{C} = \begin{bmatrix} \mathbf{I} \\ \mathbf{X} \end{bmatrix} \mathbf{W}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T\mathbf{W} \begin{bmatrix} \mathbf{I} & \mathbf{X}^T \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{X} \end{bmatrix} \mathbf{W} \begin{bmatrix} \mathbf{I} & \mathbf{X}^T \end{bmatrix}. \tag{18}$$

It follows from Equations (16) (17) (18) that $\mathbf{K} = \mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T = \mathbf{C}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T\mathbf{C}^T$.

Conversely, when $\mathbf{K} = \mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T$, it holds that $\mathrm{rank}(\mathbf{K}) \le \mathrm{rank}(\mathbf{W}^\dagger) = \mathrm{rank}(\mathbf{W})$. It follows from (15) that $\mathrm{rank}(\mathbf{K}) = \mathrm{rank}(\mathbf{W})$.

When $\mathbf{K} = \mathbf{C}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^\dagger)^T\mathbf{C}^T$, we have $\mathrm{rank}(\mathbf{K}) \le \mathrm{rank}(\mathbf{C})$. Thus there exists a matrix $\mathbf{X}$ such that

$$\begin{bmatrix} \mathbf{K}_{21}^T \\ \mathbf{K}_{22} \end{bmatrix} = \mathbf{C}\mathbf{X}^T = \begin{bmatrix} \mathbf{W}\mathbf{X}^T \\ \mathbf{K}_{21}\mathbf{X}^T \end{bmatrix},$$

and therefore $\mathbf{K}_{21} = \mathbf{X}\mathbf{W}$. Then we have that

$$\mathbf{C} = \begin{bmatrix} \mathbf{W} \\ \mathbf{K}_{21} \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ \mathbf{X} \end{bmatrix} \mathbf{W},$$

so $\mathrm{rank}(\mathbf{C}) \le \mathrm{rank}(\mathbf{W})$. Apply (15) again we have $\mathrm{rank}(\mathbf{K}) = \mathrm{rank}(\mathbf{W})$. ■

Lemma 12 $\mathbf{X}^T\mathbf{V}\mathbf{X}\big(\mathbf{X}^T\mathbf{V}\mathbf{X}\big)^\dagger\mathbf{X}^T = \mathbf{X}^T$ *for any positive definite matrix* $\mathbf{V}$.

Proof The positive definite matrix $\mathbf{V}$ have a decomposition $\mathbf{V} = \mathbf{B}^T\mathbf{B}$ for some nonsingular matrix $\mathbf{B}$. It follows that

$$\begin{aligned}\mathbf{X}^T\mathbf{V}\mathbf{X}\big(\mathbf{X}^T\mathbf{V}\mathbf{X}\big)^\dagger\mathbf{X}^T &= (\mathbf{B}\mathbf{X})^T\Big(\mathbf{B}\mathbf{X}\big((\mathbf{B}\mathbf{X})^T(\mathbf{B}\mathbf{X})\big)^\dagger\Big)(\mathbf{B}\mathbf{X})^T\mathbf{B}(\mathbf{B}^T\mathbf{B})^{-1} \\ &= (\mathbf{B}\mathbf{X})^T\big((\mathbf{B}\mathbf{X})^T\big)^\dagger(\mathbf{B}\mathbf{X})^T(\mathbf{B}^T)^{-1} = (\mathbf{B}\mathbf{X})^T(\mathbf{B}^T)^{-1} = \mathbf{X}^T.\end{aligned}$$

■

Appendix B. Proof of Theorem 2

In Section B.1 we provide several key lemmas, and then in Section B.2 we prove Theorem 2 using Lemmas 14 and 15.

B.1 Key Lemmas

Lemma 13 provides a useful tool for expanding the Moore-Penrose inverse of partitioned matrices, and the lemma will be used to prove Lemma 15 and Theorem 2.

Lemma 13 (see Ben-Israel and Greville, 2003, Page 179) *Given a matrix $\mathbf{X} \in \mathbb{R}^{m\times n}$ of rank c, let it have a nonsingular $c \times c$ submatrix $\mathbf{X}_{11}$. By rearrangement of columns and rows by permutation matrices $\mathbf{P}$ and $\mathbf{Q}$, the submatrix $\mathbf{X}_{11}$ can be bought to the top left corner of $\mathbf{X}$, that is,*

$$\mathbf{P}\mathbf{X}\mathbf{Q} = \begin{bmatrix} \mathbf{X}_{11} & \mathbf{X}_{12} \\ \mathbf{X}_{21} & \mathbf{X}_{22} \end{bmatrix}.$$

Then the Moore-Penrose inverse of $\mathbf{X}$ is

$$\mathbf{X}^\dagger = \mathbf{Q}\begin{bmatrix} \mathbf{I}_c \\ \mathbf{T}^T \end{bmatrix}\big(\mathbf{I}_c + \mathbf{T}\mathbf{T}^T\big)^{-1}\mathbf{X}_{11}^{-1}\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1}\begin{bmatrix} \mathbf{I}_c & \mathbf{S}^T \end{bmatrix}\mathbf{P},$$

where $\mathbf{T} = \mathbf{X}_{11}^{-1}\mathbf{X}_{12}$ and $\mathbf{S} = \mathbf{X}_{21}\mathbf{X}_{11}^{-1}$.

Lemmas 14 and 15 will be used to prove Theorem 2.

Lemma 14 (Wang and Zhang, 2013, Lemma 19) *Given n and k, we let $\mathbf{B}$ be an $\frac{n}{k} \times \frac{n}{k}$ matrix whose diagonal entries equal to one and off-diagonal entries equal to $\alpha \in [0, 1)$. Let $\mathbf{A}$ be the $n \times n$ block-diagonal matrix*

$$\mathbf{A} = \underbrace{\mathbf{B} \oplus \mathbf{B} \oplus \cdots \oplus \mathbf{B}}_{k \text{ blocks}}. \tag{19}$$

Let $\mathbf{A}_k$ be the best rank-k approximation to $\mathbf{A}$. Then

$$\|\mathbf{A} - \mathbf{A}_k\|_F = (1-\alpha)\sqrt{n-k}.$$

Lemma 15 *Let* $\mathbf{B}$ *be the* $n \times n$ *matrix with diagonal entries equal to one and off-diagonal entries equal to* α *and* $\tilde{\mathbf{B}}$ *be its rank* c *approximation formed by the prototype model. Then*

$$\|\mathbf{B} - \tilde{\mathbf{B}}\|_F^2 \;\geq\; (1-\alpha)^2(n-c)\left(1 + \frac{2}{c} - (1-\alpha)\frac{1+o(1)}{\alpha c n/2}\right).$$

Proof Without loss of generality, we assume the first c column of $\mathbf{B}$ are selected to construct $\mathbf{C}$. We partition $\mathbf{B}$ and $\mathbf{C}$ as:

$$\mathbf{B} = \begin{bmatrix} \mathbf{W} & \mathbf{B}_{21}^T \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{bmatrix} \qquad \text{and} \qquad \mathbf{C} = \begin{bmatrix} \mathbf{W} \\ \mathbf{B}_{21} \end{bmatrix}.$$

Here the matrix $\mathbf{W}$ can be expressed by $\mathbf{W} = (1-\alpha)\mathbf{I}_c + \alpha \mathbf{1}_c\mathbf{1}_c^T$. We apply the Sherman-Morrison-Woodbury matrix identity

$$(\mathbf{X} + \mathbf{YZR})^{-1} = \mathbf{X}^{-1} - \mathbf{X}^{-1}\mathbf{Y}(\mathbf{Z}^{-1} + \mathbf{RX}^{-1}\mathbf{Y})^{-1}\mathbf{RX}^{-1}$$

to compute $\mathbf{W}^{-1}$, and it follows that

$$\mathbf{W}^{-1} \;=\; \frac{1}{1-\alpha}\mathbf{I}_c - \frac{\alpha}{(1-\alpha)(1-\alpha+c\alpha)}\mathbf{1}_c\mathbf{1}_c^T. \tag{20}$$

We expand the Moore-Penrose inverse of $\mathbf{C}$ by Lemma 13 and obtain

$$\mathbf{C}^\dagger = \mathbf{W}^{-1}\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1} \begin{bmatrix} \mathbf{I}_c & \mathbf{S}^T \end{bmatrix}$$

where

$$\mathbf{S} = \mathbf{B}_{21}\mathbf{W}^{-1} = \frac{\alpha}{1-\alpha+c\alpha}\mathbf{1}_{n-c}\mathbf{1}_c^T.$$

It is easily verified that $\mathbf{S}^T\mathbf{S} = \big(\frac{\alpha}{1-\alpha+c\alpha}\big)^2(n-c)\mathbf{1}_c\mathbf{1}_c^T$.

Now we express the approximation by the prototype model in the partitioned form:

$$\begin{aligned}
\tilde{\mathbf{B}} \;&=\; \mathbf{C}\mathbf{C}^\dagger\mathbf{B}\big(\mathbf{C}^\dagger\big)^T\mathbf{C}^T \\
&= \begin{bmatrix} \mathbf{W} \\ \mathbf{B}_{21} \end{bmatrix} \mathbf{W}^{-1}\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1} \begin{bmatrix} \mathbf{I}_c & \mathbf{S}^T \end{bmatrix} \mathbf{B} \begin{bmatrix} \mathbf{I}_c \\ \mathbf{S} \end{bmatrix} \big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1}\mathbf{W}^{-1} \begin{bmatrix} \mathbf{W} \\ \mathbf{B}_{21} \end{bmatrix}^T \\
&= \begin{bmatrix} \big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1} \\ \mathbf{B}_{21}\mathbf{W}^{-1}\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I}_c & \mathbf{S}^T \end{bmatrix} \mathbf{B} \begin{bmatrix} \mathbf{I}_c \\ \mathbf{S} \end{bmatrix} \begin{bmatrix} \big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1} \\ \mathbf{B}_{21}\mathbf{W}^{-1}\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1} \end{bmatrix}^T.
\end{aligned} \tag{21}$$

We then compute the submatrices $\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1}$ and $\mathbf{B}_{21}\mathbf{W}^{-1}\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1}$ respectively as follows. We apply the Sherman-Morrison-Woodbury matrix identity to compute $\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1}$. We obtain

$$\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1} \;=\; \left(\mathbf{I}_c + \Big(\frac{\alpha}{1-\alpha+c\alpha}\Big)^2(n-c)\mathbf{1}_c\mathbf{1}_c^T\right)^{-1} \;=\; \mathbf{I}_c - \gamma_1\mathbf{1}_c\mathbf{1}_c^T, \tag{22}$$

where

$$\gamma_1 = \frac{n-c}{nc + \big(\frac{1-\alpha}{\alpha}\big)^2 + \frac{2(1-\alpha)c}{\alpha}}.$$

It follows from (20) and (22) that

$$\mathbf{W}^{-1}\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1} \;=\; (\gamma_2\mathbf{I}_c - \gamma_3\mathbf{1}_c\mathbf{1}_c^T)(\mathbf{I}_c - \gamma_1\mathbf{1}_c\mathbf{1}_c^T) \;=\; \gamma_2\mathbf{I}_c + (\gamma_1\gamma_3 c - \gamma_1\gamma_2 - \gamma_3)\mathbf{1}_c\mathbf{1}_c^T,$$

where

$$\gamma_2 = \frac{1}{1-\alpha} \quad \text{and} \quad \gamma_3 = \frac{\alpha}{(1-\alpha)(1-\alpha+\alpha c)}.$$

It follows that

$$\mathbf{B}_{21}\mathbf{W}^{-1}\big(\mathbf{I}_c + \mathbf{S}^T\mathbf{S}\big)^{-1} \;=\; \alpha\big(\gamma_1\gamma_3 c^2 - \gamma_3 c - \gamma_1\gamma_2 c + \gamma_2\big)\mathbf{1}_{n-c}\mathbf{1}_c^T \;\triangleq\; \gamma\mathbf{1}_{n-c}\mathbf{1}_c^T, \tag{23}$$

where

$$\gamma \;=\; \alpha\big(\gamma_1\gamma_3 c^2 - \gamma_3 c - \gamma_1\gamma_2 c + \gamma_2\big) \;=\; \frac{\alpha(\alpha c - \alpha + 1)}{2\alpha c - 2\alpha - 2\alpha^2 c + \alpha^2 + \alpha^2 cn + 1}. \tag{24}$$

Since $\mathbf{B}_{21} = \alpha\mathbf{1}_{n-c}\mathbf{1}_c^T$ and $\mathbf{B}_{22} = (1-\alpha)\mathbf{I}_{n-c} + \alpha\mathbf{1}_{n-c}\mathbf{1}_{n-c}^T$, it is easily verified that

$$\left[\begin{array}{cc} \mathbf{I}_c & \mathbf{S}^T \end{array}\right]\mathbf{B}\left[\begin{array}{c} \mathbf{I}_c \\ \mathbf{S} \end{array}\right] \;=\; \left[\begin{array}{cc} \mathbf{I}_c & \mathbf{S}^T \end{array}\right]\left[\begin{array}{cc} \mathbf{W} & \mathbf{B}_{21}^T \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{array}\right]\left[\begin{array}{c} \mathbf{I}_c \\ \mathbf{S} \end{array}\right] \;=\; (1-\alpha)\mathbf{I}_c + \lambda\mathbf{1}_c\mathbf{1}_c^T, \tag{25}$$

where

$$\lambda = \frac{\alpha(3\alpha n - \alpha c - 2\alpha + \alpha^2 c - 3\alpha^2 n + \alpha^2 + \alpha^2 n^2 + 1)}{(\alpha c - \alpha + 1)^2}$$

It follows from (21), (22), (23), and (25) that

$$\tilde{\mathbf{B}} \;=\; \left[\begin{array}{c} \mathbf{I}_c - \gamma_1\mathbf{1}_c\mathbf{1}_c^T \\ \gamma\mathbf{1}_{n-c}\mathbf{1}_c^T \end{array}\right]\Big((1-\alpha)\mathbf{I}_c + \lambda\mathbf{1}_c\mathbf{1}_c^T\Big)\left[\begin{array}{c} \mathbf{I}_c - \gamma_1\mathbf{1}_c\mathbf{1}_c^T \\ \gamma\mathbf{1}_{n-c}\mathbf{1}_c^T \end{array}\right]^T \;\triangleq\; \left[\begin{array}{cc} \tilde{\mathbf{B}}_{11} & \tilde{\mathbf{B}}_{21}^T \\ \tilde{\mathbf{B}}_{21} & \tilde{\mathbf{B}}_{22} \end{array}\right],$$

where

$$\begin{aligned}
\tilde{\mathbf{B}}_{11} &= (1-\alpha)\mathbf{I}_c + \big[(1-\gamma_1 c)(\lambda - \lambda\gamma_1 c - (1-\alpha)\gamma_1) - (1-\alpha)\gamma_1\big]\mathbf{1}_c\mathbf{1}_c^T \\
&= (1-\alpha)\mathbf{I}_c + \eta_1\mathbf{1}_c\mathbf{1}_c^T, \\
\tilde{\mathbf{B}}_{21} &= \tilde{\mathbf{A}}_{12}^T \;=\; \gamma(1-\gamma_1 c)(1-\alpha+\lambda c)\mathbf{1}_{n-c}\mathbf{1}_c^T \;=\; \eta_2\mathbf{1}_{n-c}\mathbf{1}_c^T, \\
\tilde{\mathbf{B}}_{22} &= \gamma^2 c(1-\alpha+\lambda c)\mathbf{1}_{n-c}\mathbf{1}_{n-c}^T \;=\; \eta_3\mathbf{1}_{n-c}\mathbf{1}_{n-c}^T,
\end{aligned}$$

where

$$\begin{aligned}
\eta_1 &= (1-\gamma_1 c)(\lambda - \lambda\gamma_1 c - (1-\alpha)\gamma_1) - (1-\alpha)\gamma_1, \\
\eta_2 &= \gamma(1-\gamma_1 c)(1-\alpha+\lambda c), \\
\eta_3 &= \gamma^2 c(1-\alpha+\lambda c),
\end{aligned}$$

By dealing with the four blocks of $\tilde{\mathbf{B}}$ respectively, we finally obtain that

$$\begin{aligned}
\|\mathbf{B} - \tilde{\mathbf{B}}\|_F^2 &= \|\mathbf{W} - \tilde{\mathbf{B}}_{11}\|_F^2 + 2\|\mathbf{B}_{21} - \tilde{\mathbf{B}}_{21}\|_F^2 + \|\mathbf{B}_{22} - \tilde{\mathbf{B}}_{22}\|_F^2 \\
&= c^2(\alpha - \eta_1)^2 + 2c(n-c)(\alpha - \eta_2)^2 \\
&\quad + (n-c)(n-c-1)(\alpha - \eta_3)^2 + (n-c)(1-\eta_3)^2 \\
&= (n-c)(\alpha-1)^2\left(1 + \frac{2}{c} - \big(1+o(1)\big)\frac{1-\alpha}{\alpha cn/2}\right).
\end{aligned}$$

■

B.2 Proof of the Theorem

Now we prove Theorem 2 using Lemma 15 and Lemma 14. Let $\mathbf{C}$ consist of c columns sampled from $\mathbf{A}$ and $\hat{\mathbf{C}}_i$ consist of c_i columns sampled from the i-th block diagonal matrix in $\mathbf{A}$. Without loss of generality, we assume $\hat{\mathbf{C}}_i$ consists of the first c_i columns of $\mathbf{B}$. Then the intersection matrix $\mathbf{U}$ is computed by

$$\begin{aligned}\mathbf{U} &= \mathbf{C}^\dagger\mathbf{A}(\mathbf{C}^T)^\dagger = \big[\hat{\mathbf{C}}_1 \oplus \cdots \oplus \hat{\mathbf{C}}_k\big]^\dagger\big[\mathbf{B} \oplus \cdots \oplus \mathbf{B}\big]\big[\hat{\mathbf{C}}_1^T \oplus \cdots \oplus \hat{\mathbf{C}}_k^T\big]^\dagger \\ &= \hat{\mathbf{C}}_1^\dagger\mathbf{B}(\hat{\mathbf{C}}_1^\dagger)^T \oplus \cdots \oplus \hat{\mathbf{C}}_k^\dagger\mathbf{B}(\hat{\mathbf{C}}_k^\dagger)^T.\end{aligned}$$

Let $\tilde{\mathbf{A}}$ be the approximation formed by the prototype model. Then

$$\tilde{\mathbf{A}} = \mathbf{C}\mathbf{U}\mathbf{C}^T = \hat{\mathbf{C}}_1\hat{\mathbf{C}}_1^\dagger\mathbf{B}(\hat{\mathbf{C}}_1^\dagger)^T\hat{\mathbf{C}}_1^T \oplus \cdots \oplus \hat{\mathbf{C}}_k\hat{\mathbf{C}}_k^\dagger\mathbf{B}(\hat{\mathbf{C}}_k^\dagger)^T\hat{\mathbf{C}}_k^T,$$

and thus the approximation error is

$$\begin{aligned}\|\mathbf{A} - \tilde{\mathbf{A}}\|_F^2 &= \sum_{i=1}^k \Big\|\mathbf{B} - \hat{\mathbf{C}}_i\hat{\mathbf{C}}_i^\dagger\mathbf{B}(\hat{\mathbf{C}}_i^\dagger)^T\hat{\mathbf{C}}_i^T\Big\|_F^2 \\ &\geq (1-\alpha)^2\sum_{i=1}^k(p-c_i)\Big(1 + \frac{2}{c_i} - (1-\alpha)\Big(\frac{1+o(1)}{\alpha c_i p/2}\Big)\Big) \\ &= (1-\alpha)^2\Big(\sum_{i=1}^k(p-c_i) + \sum_{i=1}^k\frac{2(p-c_i)}{c_i}\Big(1 - \frac{(1-\alpha)(1+o(1))}{\alpha p}\Big)\Big) \\ &\geq (1-\alpha)^2(n-c)\Big(1 + \frac{2k}{c}\Big(1 - \frac{k(1-\alpha)(1+o(1))}{\alpha n}\Big)\Big),\end{aligned}$$

where the former inequality follows from Lemma 15, and the latter inequality follows by minimizing over $c_1, \cdots, c_k$. Finally the theorem follows by setting $\alpha \to 1$ and applying Lemma 14.

Appendix C. Proof of Theorem 3

Proof Sampling $c_1 = 2k\epsilon^{-1}\big(1 + o(1)\big)$ columns of $\mathbf{K}$ by the near-optimal algorithm of Boutsidis et al. (2014) to form $\mathbf{C}_1 \in \mathbb{R}^{n\times c_1}$, we have that

$$\mathbb{E}\big\|\mathbf{K} - \mathcal{P}_{\mathbf{C}_1,k}(\mathbf{K})\big\|_F^2 \leq (1+\epsilon)\big\|\mathbf{K} - \mathbf{K}_k\big\|_F^2.$$

Applying Lemma 3.11 of Boutsidis and Woodruff (2014), we can find a much smaller matrix $\mathbf{U}_1 \in \mathbb{R}^{n\times k}$ with orthogonal columns in the column space of $\mathbf{C}_1$ such that

$$\big\|\mathbf{K} - \mathbf{U}_1\mathbf{U}_1^T\mathbf{K}\big\|_F^2 \leq \big\|\mathbf{K} - \mathcal{P}_{\mathbf{C}_1,k}(\mathbf{K})\big\|_F^2.$$

(We do not actually compute $\mathbf{U}_1$ because the adaptive sampling algorithm does not need to know $\mathbf{U}_1$.) Because the columns of $\mathbf{U}_1$ are all in the columns space of $\mathbf{C}_1$, we have that $\big\|\mathbf{K} - \mathbf{C}_1\mathbf{C}_1^\dagger\mathbf{K}\big\|_F^2 \leq \big\|\mathbf{K} - \mathbf{U}_1\mathbf{U}_1^T\mathbf{K}\big\|_F^2$. Combining the above inequalities we obtain

$$\mathbb{E}\big\|\mathbf{K} - \mathbf{C}_1\mathbf{C}_1^\dagger\mathbf{K}\big\|_F^2 \leq \mathbb{E}\big\|\mathbf{K} - \mathbf{U}_1\mathbf{U}_1^T\mathbf{K}\big\|_F^2 \leq \big\|\mathbf{K} - \mathcal{P}_{\mathbf{C}_1,k}(\mathbf{K})\big\|_F^2.$$

Given $\mathbf{C}_1$, we use adaptive sampling to select $c_2 = k\epsilon^{-1}$ rows of $\mathbf{K}$ to form $\mathbf{C}_2$ and denote $\mathbf{C} = [\mathbf{C}_1, \mathbf{C}_2]$. Since the columns of $\mathbf{U}_1$ are all in the columns space of $\mathbf{C}_1$, Lemma 17 of Wang and Zhang (2013) can be slightly modified to show

$$\big\|\mathbf{K} - \mathbf{C}_1\mathbf{C}_1^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\big\|_F^2 \;\leq\; \big\|\mathbf{K} - \mathbf{U}_1\mathbf{U}_1^T\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\big\|_F^2.$$

By the adaptive sampling theorem of Wang and Zhang (2013) we have

$$\begin{aligned}
\mathbb{E}\big\|\mathbf{K} - \mathbf{C}_1\mathbf{C}_1^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\big\|_F^2 &\leq \mathbb{E}\big\|\mathbf{K} - \mathbf{U}_1\mathbf{U}_1^T\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\big\|_F^2 \\
&\leq \big\|\mathbf{K} - \mathbf{U}_1\mathbf{U}_1^T\mathbf{K}\big\|_F^2 + \frac{k}{c_2}\big\|\mathbf{K} - \mathbf{K}(\mathbf{C}_1^T)^\dagger\mathbf{C}_1^T\big\|_F^2 \\
&\leq \Big(1 + \frac{k}{c_2}\Big)\big\|\mathbf{K} - \mathbf{U}_1\mathbf{U}_1^\dagger\mathbf{K}\big\|_F^2 \\
&\leq (1+\epsilon)\big\|\mathbf{K} - \mathcal{P}_{\mathbf{C}_1,k}(\mathbf{K})\big\|_F^2, \qquad (26)
\end{aligned}$$

where the expectation is taken w.r.t. $\mathbf{C}_2$, and the last inequality follows by setting $c_2 = k/\epsilon$. Here the trick is bounding $\mathbb{E}\big\|\mathbf{K} - \mathbf{U}_1\mathbf{U}_1^T\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\big\|_F^2$ rather than directly bounding $\mathbb{E}\big\|\mathbf{K} - \mathbf{C}_1\mathbf{C}_1^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\big\|_F^2$; otherwise the factor $\frac{k}{c_2}$ would be $\frac{c_1}{c_2} = \frac{2k\epsilon^{-1}(1+o(1))}{c_2}$. This is the key to the improvement. It follows that

$$\begin{aligned}
\mathbb{E}\big\|\mathbf{K} - \mathbf{C}_1\mathbf{C}_1^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\big\|_F^2 &\leq (1+\epsilon)\mathbb{E}\big\|\mathbf{K} - \mathcal{P}_{\mathbf{C}_1,k}(\mathbf{K})\big\|_F^2 \\
&\leq (1+\epsilon)^2\big\|\mathbf{K} - \mathbf{K}_k\big\|_F^2,
\end{aligned}$$

where the first expectation is taken w.r.t. $\mathbf{C}_1$ and $\mathbf{C}_2$, and the second expectation is taken w.r.t. $\mathbf{C}_1$. Applying Lemma 17 of Wang and Zhang (2013) again, we obtain

$$\mathbb{E}\big\|\mathbf{K} - \mathbf{C}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\big\|_F^2 \;\leq\; \mathbb{E}\big\|\mathbf{K} - \mathbf{C}_1\mathbf{C}_1^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\big\|_F^2 \;\leq\; (1+\epsilon)^2\big\|\mathbf{K} - \mathbf{K}_k\big\|_F^2.$$

Hence totally $c = c_1 + c_2 = 3k\epsilon^{-1}\big(1 + o(1)\big)$ columns suffice. ■

Appendix D. Proof of Theorem 4

In this section we first provide a constant factor bound of the uniform sampling, and then prove Theorem 4 in the subsequent subsections.

D.1 Tools for Analyzing Uniform Sampling

This subsection provides several useful tools for analyzing column sampling. The matrix $\mathbf{S} \in \mathbb{R}^{n\times s}$ is a column selection matrix if each column has exactly one nonzero entry; let (i_j, j) be the position of the nonzero entry in the j-th column. Let us add randomness to column selection. Suppose we are given the sampling probabilities $p_1, \cdots, p_n \in (0,1)$ and $\sum_i p_i = 1$. In each round we pick one element in $[n]$ such that the i-th element is sampled with probability p_i. We repeat the procedure s times, either with or without replacement, and let $i_1, \cdots, i_s$ be the selected indices. For $j = 1$ to s, we set

$$S_{i_j,j} = \frac{1}{\sqrt{sp_{i_j}}}. \qquad (27)$$

The following lemma shows an important property of arbitrary column sampling. The proof mirrors the leverage score sampling bound in (Woodruff, 2014).

Lemma 16 *Let $\mathbf{U} \in \mathbb{R}^{n\times k}$ be any fixed matrix with orthonormal columns. The column selection matrix $\mathbf{S} \in \mathbb{R}^{n\times s}$ samples s columns according to arbitrary probabilities $p_1, p_2, \cdots, p_n$. Assume that*

$$\max_{i\in[n]} \frac{\|\mathbf{u}_{i:}\|_2^2}{p_i} \;\leq\; \alpha$$

and $\alpha \geq 1$. When $s \;\geq\; \alpha\frac{6+2\eta}{3\eta^2}\log(k/\delta)$, it holds that

$$\mathbb{P}\Big\{\big\|\mathbf{I}_k - \mathbf{U}^T\mathbf{S}\mathbf{S}^T\mathbf{U}\big\|_2 \;\geq\; \eta\Big\} \;\leq\; \delta.$$

Proof We can express $\mathbf{I}_k = \mathbf{U}^T\mathbf{U}$ as the sum of size $k\times k$ and rank one matrices:

$$\mathbf{I}_k \;\;=\;\; \mathbf{U}^T\mathbf{U} \;=\; \sum_{i=1}^{n} \mathbf{u}_{i:}^T\mathbf{u}_{i:}$$

where $\mathbf{u}_{i:} \in \mathbb{R}^{1\times k}$ is the i-th row of $\mathbf{U}$. The approximate matrix product can be expressed as

$$\mathbf{U}^T\mathbf{S}\mathbf{S}^T\mathbf{U} \;\;=\;\; \sum_{j=1}^{s} S_{i_j,j}^2\mathbf{u}_{i_j:}^T\mathbf{u}_{i_j:} \;=\; \sum_{j=1}^{s}\frac{1}{sp_{i_j}}\mathbf{u}_{i_j:}^T\mathbf{u}_{i_j:}.$$

We define the symmetric random matrices

$$\mathbf{Z}_{i_j} \;=\; \frac{1}{sp_{i_j}}\mathbf{u}_{i_j:}^T\mathbf{u}_{i_j:} - \frac{1}{s}\mathbf{I}_k$$

for $j = 1$ to s. Whatever the sampling distribution is, it always holds that

$$\mathbb{E}\mathbf{Z}_{i_j} \;=\; \sum_{q=1}^{n} p_q\Big(\frac{1}{\sqrt{sp_q}}\Big)^2\mathbf{u}_{q:}^T\mathbf{u}_{q:} - \sum_{q=1}^{n} p_q\frac{1}{s}\mathbf{I}_k \;=\; \frac{1}{s}\mathbf{I}_k - \frac{1}{s}\mathbf{I}_k \;=\; \mathbf{0}. \tag{28}$$

Thus $\mathbf{Z}_{i_j}$ has zero mean. Let $\mathcal{Z}$ be the set

$$\mathcal{Z} \;=\; \Big\{\frac{1}{s}\mathbf{I}_k - \frac{1}{sp_1}\mathbf{u}_{1:}^T\mathbf{u}_{1:},\; \cdots,\; \frac{1}{s}\mathbf{I}_k - \frac{1}{sp_n}\mathbf{u}_{n:}^T\mathbf{u}_{n:}\Big\}.$$

Clearly, $\mathbf{Z}_{i_1}, \cdots, \mathbf{Z}_{i_s}$ are sampled from $\mathcal{Z}$, and

$$\mathbf{U}^T\mathbf{S}\mathbf{S}^T\mathbf{U} - \mathbf{I}_k = \sum_{j=1}^{s}\mathbf{Z}_{i_j}.$$

Therefore we can bound its spectral norm using the matrix Bernstein. Elementary proof of the matrix Bernstein can be found in Tropp (2015).

Lemma 17 (Matrix Bernstein) *Consider a finite sequence $\{\mathbf{Z}_i\}$ of independent, random, Hermitian matrices with dimension k. Assume that*

$$\mathbb{E}\mathbf{Z}_i = \mathbf{0} \quad \text{and} \quad \max_i \|\mathbf{Z}_i\|_2 \le L$$

for each index i. Introduce the random matrix $\mathbf{Y} = \sum_i \mathbf{Z}_i$. Let $v(\mathbf{Y})$ be the matrix variance statistics of the sum:

$$v(\mathbf{Y}) \;=\; \left\|\mathbb{E}\mathbf{Y}^2\right\|_2 \;=\; \left\|\sum_i \mathbb{E}\mathbf{Z}_i^2\right\|_2.$$

Then

$$\mathbb{P}\{\lambda_{\max}(\mathbf{Y}) \ge \eta\} \;\le\; k \cdot \exp\left(\frac{-\eta^2/2}{v(\mathbf{Y}) + L\eta/3}\right).$$

To apply the matrix Bernstein, it remains to bound the variance of $\sum_{j=1}^s \mathbf{Z}_{i_j}$ and to bound $L = \max_{j\in[s]} \|\mathbf{Z}_{i_j}\|_2$. We have that

$$\begin{aligned}
\mathbb{E}\mathbf{Z}_{i_j}^2 &= \mathbb{E}\Big(\frac{1}{sp_{i_j}}\mathbf{u}_{i_j:}^T\mathbf{u}_{i_j:} - \frac{1}{s}\mathbf{I}_k\Big)^2 \\
&= \frac{1}{s^2}\left[\mathbb{E}\Big(\frac{1}{p_{i_j}}\mathbf{u}_{i_j:}^T\mathbf{u}_{i_j:}\Big)^2 - 2\mathbb{E}\Big(\frac{1}{p_{i_j}}\mathbf{u}_{i_j:}^T\mathbf{u}_{i_j:}\Big) + \mathbf{I}_k\right] \\
&= \frac{1}{s^2}\left[\sum_{q=1}^n \Big(p_q\Big(\frac{1}{p_q}\Big)^2\mathbf{u}_{q:}^T\mathbf{u}_{q:}\mathbf{u}_{q:}^T\mathbf{u}_{q:}\Big) - 2\mathbf{I}_k + \mathbf{I}_k\right] \\
&= -\frac{1}{s^2}\mathbf{I}_k + \frac{1}{s^2}\sum_{q=1}^n \frac{\|\mathbf{u}_{q:}\|_2^2}{p_q}\mathbf{u}_{q:}^T\mathbf{u}_{q:}.
\end{aligned}$$

The variance is defined by

$$v \;\triangleq\; \left\|\sum_{j=1}^s \mathbb{E}\mathbf{Z}_{i_j}^2\right\|_2 \;=\; \frac{1}{s}\left\|-\mathbf{I}_k + \sum_{q=1}^n \frac{\|\mathbf{u}_{q:}\|_2^2}{p_q}\mathbf{u}_{q:}^T\mathbf{u}_{q:}\right\|_2 \;\le\; \frac{1}{s}(\alpha - 1).$$

Here the inequality is due to the definition of α and

$$-\mathbf{I}_k + \sum_{q=1}^n \frac{\|\mathbf{u}_{q:}\|_2^2}{p_q}\mathbf{u}_{q:}^T\mathbf{u}_{q:} \;\preceq\; -\mathbf{I}_k + \sum_{q=1}^n \alpha\mathbf{u}_{q:}^T\mathbf{u}_{q:} \;=\; (-1+\alpha)\mathbf{I}_k.$$

In addition, $L = \max_{j\in[s]} \|\mathbf{Z}_{i_j}\|_2$ can be bounded by

$$\begin{aligned}
L \;&=\; \max_{j\in[s]} \|\mathbf{Z}_{i_j}\|_2 \;=\; \max_{j\in[s]} \left\|\frac{1}{s}\mathbf{I}_k - \frac{1}{sp_{i_j}}\mathbf{u}_{i_j:}^T\mathbf{u}_{i_j:}\right\|_2 \;\le\; \max_{i\in[n]} \left\|\frac{1}{s}\mathbf{I}_k - \frac{1}{sp_i}\mathbf{u}_{i:}^T\mathbf{u}_{i:}\right\|_2 \\
&\le\; \frac{1}{s} + \max_{i\in[n]} \left\|\frac{1}{sp_i}\mathbf{u}_{i:}^T\mathbf{u}_{i:}\right\|_2 \;=\; \frac{1}{s} + \max_{i\in[n]} \frac{\|\mathbf{u}_{i:}\|_2^2}{sp_i} \;\le\; \frac{1}{s}(\alpha+1).
\end{aligned}$$

Finally, the lemma follows by plugging v and L in the matrix Bernstein. ∎

Theorem 18 shows an important property of uniform sampling. It shows that when the number of sampled columns is large enough, all the singular values of $\mathbf{U}^T\mathbf{S}\mathbf{S}^T\mathbf{U}$ are within $1 \pm \eta$ with high probability.

Theorem 18 *Let $\mathbf{U} \in \mathbb{R}^{n\times k}$ be any fixed matrix with orthonormal columns and $\mu(\mathbf{U})$ be its row coherence. Let $\mathbf{S} \in \mathbb{R}^{n\times s}$ be an uniformly sampling matrix. When $s \geq \mu(\mathbf{U})k\frac{6+2\eta}{3\eta^2}\log(k/\delta)$, it hold that*

$$\mathbb{P}\Big\{\big\|\mathbf{U}^T\mathbf{S}\mathbf{S}^T\mathbf{U} - \mathbf{I}_k\big\|_2 \;\geq\; \eta\Big\} \;\leq\; \delta.$$

Proof Since $\frac{\|\mathbf{u}_{q:}\|_2^2}{p_q} = n\|\mathbf{u}_{q:}\|_2^2 \leq k\mu(\mathbf{U})$ for all $q \in [n]$, the theorem follows from Lemma 16. ■

The following lemma was established by Drineas et al. (2008). It can be proved by writing the squared Frobenius norm as the sum of scalars and then taking expectation.

Lemma 19 *Let $\mathbf{A} \in \mathbb{R}^{n\times k}$ and $\mathbf{B} \in \mathbb{R}^{n\times d}$ be any fixed matrices and $\mathbf{S} \in \mathbb{R}^{n\times s}$ be the column sampling matrix defined in (27). Assume that the columns are selected randomly and pairwisely independently. Then*

$$\mathbb{E}\Big\|\mathbf{A}^T\mathbf{B} - \mathbf{A}^T\mathbf{S}\mathbf{S}^T\mathbf{B}\Big\|_F^2 \;\leq\; \frac{1}{s}\sum_{i=1}^{n}\frac{1}{p_i}\big\|\mathbf{a}_{i:}\big\|_2^2\big\|\mathbf{b}_{i:}\big\|_2^2.$$

Theorem 20 follows from Lemma 19 and shows another important property of uniform sampling.

Theorem 20 *Let $\mathbf{U} \in \mathbb{R}^{n\times k}$ be any fixed matrix with orthonormal columns and $\mathbf{B} \in \mathbb{R}^{n\times d}$ be any fixed matrix. Use the uniform sampling matrix $\mathbf{S} \in \mathbb{R}^{n\times s}$ and set $s \geq \frac{k\mu(\mathbf{U})}{\epsilon\delta}$. Then it holds that*

$$\mathbb{P}\Big\{\big\|\mathbf{U}^T\mathbf{B} - \mathbf{U}^T\mathbf{S}^T\mathbf{S}\mathbf{B}\big\|_F^2 \;\geq\; \epsilon\|\mathbf{B}\|_F^2\Big\} \;\leq\; \delta.$$

Proof We have shown that $\|\mathbf{u}_{q:}\|_2^2/p_q \leq k\mu(\mathbf{U})$ for all $q = 1$ to n. It follows from Lemma 19 that

$$\mathbb{E}\Big\|\mathbf{U}^T\mathbf{B} - \mathbf{U}^T\mathbf{S}\mathbf{S}^T\mathbf{B}\Big\|_F^2 \;\leq\; \frac{1}{s}\sum_{i=1}^{n}k\mu(\mathbf{U})\big\|\mathbf{b}_{i:}\big\|_2^2 \;=\; \frac{k\mu(\mathbf{U})}{s}\|\mathbf{B}\|_F^2.$$

The theorem follows from the setting of s and the Markov's inequality. ■

The following lemma is very useful in analyzing randomized SVD.

Lemma 21 *Let $\mathbf{M} \in \mathbb{R}^{m\times n}$ be any matrix. We decompose $\mathbf{M}$ by $\mathbf{M} = \mathbf{M}_1 + \mathbf{M}_2$ such that $\mathrm{rank}(\mathbf{M}_1) = k$. Let the right singular vectors of $\mathbf{M}_1$ be $\mathbf{V}_1 \in \mathbb{R}^{n\times k}$. Let $\mathbf{S} \in \mathbb{R}^{n\times c}$ be any matrix such that $\mathrm{rank}(\mathbf{V}_1^T\mathbf{S}) = k$ and let $\mathbf{C} = \mathbf{M}\mathbf{S} \in \mathbb{R}^{m\times c}$. Then*

$$\big\|\mathbf{M} - \mathbf{C}(\mathbf{V}_1^T\mathbf{S})^\dagger\mathbf{V}_1^T\big\|_\xi^2 \;\leq\; \big\|\mathbf{M}_2\big\|_\xi^2 + \sigma_{\min}^{-2}(\mathbf{V}_1^T\mathbf{S}\mathbf{S}^T\mathbf{V}_1)\,\big\|\mathbf{M}_2\mathbf{S}\mathbf{S}^T\mathbf{V}_1\big\|_\xi^2$$

for $\xi = 2$ or F.

Proof Boutsidis et al. (2014) showed that

$$\left\|\mathbf{M}-\mathbf{C}(\mathbf{V}_1^T\mathbf{S})^\dagger\mathbf{V}_1^T\right\|_\xi^2 \;\leq\; \left\|\mathbf{M}_2\right\|_\xi^2+\left\|\mathbf{M}_2\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\right\|_\xi^2.$$

Since $\mathbf{Y}^T(\mathbf{Y}\mathbf{Y}^T)^\dagger=\mathbf{Y}^\dagger$ for any matrix $\mathbf{Y}$, it follows that

$$\begin{aligned}&\left\|\mathbf{M}_2\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^\dagger\right\|_\xi^2 \;=\; \left\|\mathbf{M}_2\mathbf{S}(\mathbf{V}_1^T\mathbf{S})^T(\mathbf{V}_1^T\mathbf{S}\mathbf{S}^T\mathbf{V}_1)^\dagger\right\|_\xi^2\\ &\leq\; \left\|\mathbf{M}_2\mathbf{S}\mathbf{S}^T\mathbf{V}_1\right\|_\xi^2\left\|(\mathbf{V}_1^T\mathbf{S}\mathbf{S}^T\mathbf{V}_1)^\dagger\right\|_2^2 \;=\; \left\|\mathbf{M}_2\mathbf{S}\mathbf{S}^T\mathbf{V}_1^T\right\|_\xi^2\sigma_{\min}^{-2}(\mathbf{V}_1^T\mathbf{S}\mathbf{S}^T\mathbf{V}_1),\end{aligned}$$

by which the lemma follows. ■

D.2 Uniform Sampling Bound

Theorem 22 shows that uniform sampling can be applied to randomized SVD. Specifically, if $\mathbf{C}$ consists of $c=\mathcal{O}(k\mu_k/\epsilon+k\mu_k\log k)$ uniformly sampled columns of $\mathbf{A}$, then $\|\mathbf{A}-\mathcal{P}_{\mathbf{C},k}(\mathbf{A})\|_F^2\leq(1+\epsilon)\|\mathbf{A}-\mathbf{A}_k\|_F^2$ holds with high probability, where μ_k is the column coherence of $\mathbf{A}_k$.

Theorem 22 *Let $\mathbf{A}\in\mathbb{R}^{m\times n}$ be any fixed matrix, k ($\ll m,n$) be the target rank, and μ_k be the column coherence of $\mathbf{A}_k$. Let $\mathbf{S}\in\mathbb{R}^{m\times c}$ be a uniform sampling matrix and $\mathbf{C}=\mathbf{A}\mathbf{S}\in\mathbb{R}^{m\times c}$. When $c\geq\mu_k k\cdot\max\{20\log(20k),\,45/\epsilon\}$,*

$$\min_{\mathrm{rank}(\mathbf{X})\leq k}\left\|\mathbf{A}-\mathbf{C}\mathbf{X}\right\|_F^2 \;\leq\; \left(1+\epsilon\right)\left\|\mathbf{A}-\mathbf{A}_k\right\|_F^2.$$

holds with probability at least 0.9.

Proof Let $\mathbf{V}_k\in\mathbb{R}^{n\times k}$ contain the top k right singular vectors of $\mathbf{A}$. Obviously μ_k is the row coherence of $\mathbf{V}_k$. We apply Theorem 18 with $\delta=0.05$, $\eta=1/3$, $s=20\mu_k k\log(20k)$ and obtain

$$\mathbb{P}\Big\{\sigma_{\min}(\mathbf{V}_k^T\mathbf{S}\mathbf{S}^T\mathbf{V}_k)\;\leq\;2/3\Big\}\;\leq\;0.05.$$

We then apply Theorem 20 with $\delta=0.05$ and $s=45\mu_k k/\epsilon$ and obtain

$$\mathbb{P}\Big\{\left\|\mathbf{U}^T\mathbf{S}^T\mathbf{S}(\mathbf{A}-\mathbf{A}_k)\right\|_F^2\;\geq\;\frac{4}{9}\epsilon\|\mathbf{A}-\mathbf{A}_k\|_F^2\Big\}\;\leq\;0.05.$$

It follows from Lemma 21 that

$$\begin{aligned}\min_{\mathrm{rank}(\mathbf{X})\leq k}\|\mathbf{A}-\mathbf{C}\mathbf{X}\|_F^2 \;&\leq\; \|\mathbf{A}-\mathbf{A}_k\|_F^2+\sigma_{\min}^{-2}(\mathbf{V}_k^T\mathbf{S}\mathbf{S}^T\mathbf{V}_k)\left\|\mathbf{U}^T\mathbf{S}^T\mathbf{S}(\mathbf{A}-\mathbf{A}_k)\right\|_F^2\\ &\leq\; (1+\epsilon)\|\mathbf{A}-\mathbf{A}_k\|_F^2,\end{aligned}$$

where the latter inequality holds with probability at least 0.9 (due to the union bound). ■

D.3 The Uniform+Adaptive Column Selection Algorithm

In fact, running adaptive sampling only once yields a column sampling algorithm with $1+\epsilon$ bound for any $m \times n$ matrix. We call it the uniform+adaptive column selection algorithm, which is a part of the uniform+adaptive^2 algorithm. Though the algorithm is irrelevant to this work, we describe it in the following for it is of independent interest.

Theorem 23 (The Uniform+Adaptive Algorithm) *Given an $m \times n$ matrix $\mathbf{A}$, we sample $c_1 = 20\mu_k k \log(20k)$ columns by uniform sampling to form $\mathbf{C}_1$ and sample additional $c_2 = 17.5k/\epsilon$ columns by adaptive sampling to form $\mathbf{C}_2$. Let $\mathbf{C} = [\mathbf{C}_1, \mathbf{C}_2]$. Then the inequality*

$$\min_{\mathrm{rank}(\mathbf{X}) \le k} \big\|\mathbf{A} - \mathbf{C}\mathbf{X}\big\|_F^2 \;\le\; (1+\epsilon)\big\|\mathbf{A} - \mathbf{A}_k\big\|_F^2$$

holds with probability at least 0.8.

Proof We apply Theorem 22 with $\epsilon = 0.75$ and $c_1 = 20\mu_k k \log(20k)$ and obtain that

$$\|\mathbf{A} - \mathbf{C}_1\mathbf{C}_1^\dagger \mathbf{A}\|_F^2 \;\le\; 1.75\|\mathbf{A} - \mathbf{A}_k\|_F^2$$

holds with probability at least 0.9.

Deshpande et al. (2006) showed that

$$\mathbb{E}\Big[\min_{\mathrm{rank}(\mathbf{X}) \le k} \big\|\mathbf{A} - \mathbf{C}\mathbf{X}\big\|_F^2 - \|\mathbf{A} - \mathbf{A}_k\|_F^2\Big] \;\le\; \frac{k}{c_2}\|\mathbf{A} - \mathbf{C}_1\mathbf{C}_1^\dagger \mathbf{A}\|_F^2.$$

It follows from Markov's inequality that

$$\min_{\mathrm{rank}(\mathbf{X}) \le k} \big\|\mathbf{A} - \mathbf{C}\mathbf{X}\big\|_F^2 - \|\mathbf{A} - \mathbf{A}_k\|_F^2 \;\le\; \frac{k}{\delta_2 c_2}\|\mathbf{A} - \mathbf{C}_1\mathbf{C}_1^\dagger \mathbf{A}\|_F^2$$

holds with probability at least $1 - \delta_2$. We let $\delta_2 = 0.1$ and $c_2 = 17.5k/\epsilon$, and it follows that

$$\min_{\mathrm{rank}(\mathbf{X}) \le k} \big\|\mathbf{A} - \mathbf{C}\mathbf{X}\big\|_F^2 \;\le\; (1+\epsilon)\|\mathbf{A} - \mathbf{A}_k\|_F^2$$

holds with probability at least 0.8. ■

D.4 Proof of Theorem 4

We sample $c_1 = 20\mu_k k \log(20k)$ columns by uniform sampling to form $\mathbf{C}_1$ and sample additional $c_2 = 17.5k/\epsilon$ columns by adaptive sampling to form $\mathbf{C}_2$. Let $\hat{\mathbf{C}} = [\mathbf{C}_1, \mathbf{C}_2]$. It follows from Theorem 23 that

$$\big\|\mathbf{K} - \mathcal{P}_{\hat{\mathbf{C}},k}(\mathbf{K})\big\|_F^2 \;\le\; (1+\epsilon)\big\|\mathbf{K} - \mathbf{K}_k\big\|_F^2$$

holds with probability at least 0.8.

Let $\mathbf{C} = [\hat{\mathbf{C}}, \mathbf{C}_3]$ where $\mathbf{C}_3$ consists of c_3 columns of $\mathbf{K}$ chosen by adaptive sampling. Then

$$\begin{aligned}\mathbb{E}\|\mathbf{K} - \mathbf{C}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\|_F^2 &\leq \mathbb{E}\|\mathbf{K} - \hat{\mathbf{C}}\hat{\mathbf{C}}^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\|_F^2 \\ &\leq \Big(1 + \frac{k}{c_3}\Big)\|\mathbf{K} - \mathcal{P}_{\hat{\mathbf{C}},k}(\mathbf{K})\|_F^2,\end{aligned}$$

where the former inequality follows from Lemma 17 of Wang and Zhang (2013), and the latter inequality follows from (26). It follows from Markov's inequality that

$$\|\mathbf{K} - \mathbf{C}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\|_F^2 - \|\mathbf{K} - \mathcal{P}_{\hat{\mathbf{C}},k}(\mathbf{K})\|_F^2 \leq \frac{k}{\delta_3 c_3}\|\mathbf{K} - \mathcal{P}_{\hat{\mathbf{C}},k}(\mathbf{K})\|_F^2$$

holds with probability at least $1 - \delta_3$. We set $\delta_3 = 0.1$ and $c_3 = 10k/\epsilon$, Then

$$\|\mathbf{K} - \mathbf{C}\mathbf{C}^\dagger\mathbf{K}(\mathbf{C}^T)^\dagger\mathbf{C}^T\|_F^2 \leq (1+\epsilon)\|\mathbf{K} - \mathcal{P}_{\hat{\mathbf{C}},k}(\mathbf{K})\|_F^2 \leq (1+\epsilon)^2\|\mathbf{K} - \mathbf{K}_k\|_F^2,$$

where the former inequality holds with probability $1 - \delta_3$, and the latter inequality holds with probability $0.8 - \delta_3 = 0.7$.

Appendix E. Proof of Theorem 6

In Section E.1 we derive the solution to the optimization problem (14). In Section E.2 we prove that the solutions are global optimum. In Section E.3 we prove that the resulting solution is positive (semi)definite when $\mathbf{K}$ is positive (semi)definite.

E.1 Solution to the Optimization Problem (14)

We denote the objective function of the optimization problem (14) by

$$f(\mathbf{U}, \delta) = \|\mathbf{K} - \bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T - \delta\mathbf{I}_n\|_F^2.$$

We take the derivative of $f(\mathbf{U}, \delta)$ w.r.t. $\mathbf{U}$ to be zero

$$\begin{aligned}\frac{\partial f(\mathbf{U},\delta)}{\partial \mathbf{U}} &= \frac{\partial}{\partial \mathbf{U}}\mathrm{tr}(\bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T\bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T - 2\mathbf{K}\bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T + 2\delta\bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T) \\ &= 2\bar{\mathbf{C}}^T\bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T\bar{\mathbf{C}} - 2\bar{\mathbf{C}}^T\mathbf{K}\bar{\mathbf{C}} + 2\delta\bar{\mathbf{C}}^T\bar{\mathbf{C}} = \mathbf{0},\end{aligned}$$

and obtain the solution

$$\begin{aligned}\mathbf{U}^{\mathrm{ss}} &= (\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger(\bar{\mathbf{C}}^T\mathbf{K}\bar{\mathbf{C}} - \delta^{\mathrm{ss}}\bar{\mathbf{C}}^T\bar{\mathbf{C}})(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger \\ &= (\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger\bar{\mathbf{C}}^T\mathbf{K}\bar{\mathbf{C}}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger - \delta^{\mathrm{ss}}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger\bar{\mathbf{C}}^T\bar{\mathbf{C}}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger \\ &= \bar{\mathbf{C}}^\dagger\mathbf{K}(\bar{\mathbf{C}}^\dagger)^T - \delta^{\mathrm{ss}}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger.\end{aligned}$$

Similarly, we take the derivative of $f(\mathbf{U}, \delta)$ w.r.t. δ to be zero

$$\frac{\partial f(\mathbf{U},\delta)}{\partial \delta} = \frac{\partial}{\partial \delta}\mathrm{tr}(\delta^2\mathbf{I}_n - 2\delta\mathbf{K} + 2\delta\bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T) = 2n\delta - 2\mathrm{tr}(\mathbf{K}) + 2\mathrm{tr}(\bar{\mathbf{C}}\mathbf{U}\bar{\mathbf{C}}^T) = 0,$$

and it follows that

$$\begin{aligned}
\delta^{\text{ss}} &= \frac{1}{n}\Big(\text{tr}(\mathbf{K}) - \text{tr}(\bar{\mathbf{C}}\mathbf{U}^{\text{ss}}\bar{\mathbf{C}}^T)\Big) \\
&= \frac{1}{n}\Big(\text{tr}(\mathbf{K}) - \text{tr}\Big(\bar{\mathbf{C}}\bar{\mathbf{C}}^\dagger\mathbf{K}(\bar{\mathbf{C}}^\dagger)^T\bar{\mathbf{C}}^T\Big) + \delta^{\text{ss}}\text{tr}\Big(\bar{\mathbf{C}}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger\bar{\mathbf{C}}^T\Big)\Big) \\
&= \frac{1}{n}\Big(\text{tr}(\mathbf{K}) - \text{tr}\Big(\bar{\mathbf{C}}^T\bar{\mathbf{C}}\bar{\mathbf{C}}^\dagger\mathbf{K}(\bar{\mathbf{C}}^\dagger)^T\Big) + \delta^{\text{ss}}\text{tr}\Big(\bar{\mathbf{C}}\bar{\mathbf{C}}^\dagger\Big)\Big) \\
&= \frac{1}{n}\Big(\text{tr}(\mathbf{K}) - \text{tr}\Big(\bar{\mathbf{C}}^T\mathbf{K}(\bar{\mathbf{C}}^\dagger)^T\Big) + \delta^{\text{ss}}\text{rank}(\bar{\mathbf{C}})\Big),
\end{aligned}$$

and thus

$$\delta^{\text{ss}} = \frac{1}{n - \text{rank}(\bar{\mathbf{C}})}\Big(\text{tr}(\mathbf{K}) - \text{tr}\big(\bar{\mathbf{C}}^\dagger\mathbf{K}\bar{\mathbf{C}}\big)\Big).$$

E.2 Proof of Optimality

The Hessian matrix of $f(\mathbf{U}, \delta)$ w.r.t. $(\mathbf{U}, \delta)$ is

$$\mathbf{H} = \begin{bmatrix} \dfrac{\partial^2 f(\mathbf{U},\delta)}{\partial \mathsf{vec}(\mathbf{U})\partial \mathsf{vec}(\mathbf{U})^T} & \dfrac{\partial^2 f(\mathbf{U},\delta)}{\partial \mathsf{vec}(\mathbf{U})\partial\delta} \\ \dfrac{\partial^2 f(\mathbf{U},\delta)}{\partial\delta\partial \mathsf{vec}(\mathbf{U})^T} & \dfrac{\partial^2 f(\mathbf{U},\delta)}{\partial\delta^2} \end{bmatrix} = 2\begin{bmatrix} (\bar{\mathbf{C}}^T\bar{\mathbf{C}}) \otimes (\bar{\mathbf{C}}^T\bar{\mathbf{C}}) & \mathsf{vec}(\bar{\mathbf{C}}^T\bar{\mathbf{C}}) \\ \mathsf{vec}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^T & n \end{bmatrix}.$$

Here $\otimes$ denotes the Kronecker product, and $\mathsf{vec}(\mathbf{A})$ denotes the vectorization of the matrix $\mathbf{A}$ formed by stacking the columns of $\mathbf{A}$ into a single column vector. For any $\mathbf{X} \in \mathbb{R}^{c\times c}$ and $b \in \mathbb{R}$, we let

$$\begin{aligned}
q(\mathbf{X}, b) &= \begin{bmatrix}\mathsf{vec}(\mathbf{X})^T & b\end{bmatrix}\mathbf{H}\begin{bmatrix}\mathsf{vec}(\mathbf{X}) \\ b\end{bmatrix} \\
&= \mathsf{vec}(\mathbf{X})^T\big((\bar{\mathbf{C}}^T\bar{\mathbf{C}}) \otimes (\bar{\mathbf{C}}^T\bar{\mathbf{C}})\big)\mathsf{vec}(\mathbf{X}) + 2b\,\mathsf{vec}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^T\mathsf{vec}(\mathbf{X}) + nb^2 \\
&= \mathsf{vec}(\mathbf{X})^T\mathsf{vec}\big((\bar{\mathbf{C}}^T\bar{\mathbf{C}})\mathbf{X}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})\big) + 2b\,\mathsf{vec}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^T\mathsf{vec}(\mathbf{X}) + nb^2 \\
&= \text{tr}(\mathbf{X}^T\bar{\mathbf{C}}^T\bar{\mathbf{C}}\mathbf{X}\bar{\mathbf{C}}^T\bar{\mathbf{C}}) + 2b\,\text{tr}(\bar{\mathbf{C}}^T\bar{\mathbf{C}}\mathbf{X}) + nb^2.
\end{aligned}$$

Let $\bar{\mathbf{C}}\mathbf{X}\bar{\mathbf{C}}^T = \mathbf{Y} \in \mathbb{R}^{n\times n}$. Then

$$\begin{aligned}
q(\mathbf{X}, b) &= \text{tr}(\bar{\mathbf{C}}\mathbf{X}^T\bar{\mathbf{C}}^T\bar{\mathbf{C}}\mathbf{X}\bar{\mathbf{C}}^T) + 2b\,\text{tr}(\bar{\mathbf{C}}\mathbf{X}\bar{\mathbf{C}}^T) + nb^2 \\
&= \text{tr}(\mathbf{Y}^T\mathbf{Y}) + 2b\,\text{tr}(\mathbf{Y}) + nb^2 \\
&= \sum_{i=1}^n\sum_{j=1}^n y_{ij}^2 + 2b\sum_{l=1}^n y_{ll} + nb^2 \\
&= \sum_{i\neq j} y_{ij}^2 + \sum_{l=1}^n (y_{ll} + b)^2 \\
&\geq 0,
\end{aligned}$$

which shows that the Hessian matrix $\mathbf{H}$ is SPSD. Hence $f(\mathbf{U}^{\text{ss}}, \delta^{\text{ss}})$ is the global minimum of f.

E.3 Proof of Positive (Semi)Definite

We denote the thin SVD of $\bar{\mathbf{C}}$ by $\bar{\mathbf{C}} = \mathbf{U}_{\bar{\mathbf{C}}}\mathbf{\Sigma}_{\bar{\mathbf{C}}}\mathbf{V}_{\bar{\mathbf{C}}}^T$ and let $\mathbf{U}_{\bar{\mathbf{C}}}^\perp$ be the orthogonal complement of $\mathbf{U}_{\bar{\mathbf{C}}}$. The approximation is

$$\begin{aligned}\tilde{\mathbf{K}} &= \bar{\mathbf{C}}\mathbf{U}^{\text{ss}}\bar{\mathbf{C}}^T + \delta^{\text{ss}}\mathbf{I}_n = \bar{\mathbf{C}}\big(\bar{\mathbf{C}}^\dagger\mathbf{K}(\bar{\mathbf{C}}^\dagger)^T - \delta^{\text{ss}}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger\big)\bar{\mathbf{C}}^T + \delta^{\text{ss}}\mathbf{I}_n \\ &= \bar{\mathbf{C}}\big(\bar{\mathbf{C}}^\dagger\mathbf{K}(\bar{\mathbf{C}}^\dagger)^T\big)\bar{\mathbf{C}}^T + \delta^{\text{ss}}\big(\mathbf{I}_n - \bar{\mathbf{C}}(\bar{\mathbf{C}}^T\bar{\mathbf{C}})^\dagger\bar{\mathbf{C}}^T\big) \\ &= \mathbf{U}_{\bar{\mathbf{C}}}\mathbf{U}_{\bar{\mathbf{C}}}^T\mathbf{K}\mathbf{U}_{\bar{\mathbf{C}}}\mathbf{U}_{\bar{\mathbf{C}}}^T + \delta^{\text{ss}}\big(\mathbf{I}_n - \mathbf{U}_{\bar{\mathbf{C}}}\mathbf{U}_{\bar{\mathbf{C}}}^T\big) \\ &= \mathbf{U}_{\bar{\mathbf{C}}}\mathbf{U}_{\bar{\mathbf{C}}}^T\mathbf{K}\mathbf{U}_{\bar{\mathbf{C}}}\mathbf{U}_{\bar{\mathbf{C}}}^T + \delta^{\text{ss}}\mathbf{U}_{\bar{\mathbf{C}}}^\perp(\mathbf{U}_{\bar{\mathbf{C}}}^\perp)^T. \end{aligned} \tag{29}$$

The first term is SPSD because $\mathbf{K}$ is SPSD. The second term is SPSD if δ^{ss} is nonnegative. We have

$$\begin{aligned}\delta^{\text{ss}} &= \operatorname{tr}(\mathbf{K}) - \operatorname{tr}\big(\bar{\mathbf{C}}^\dagger\mathbf{K}\bar{\mathbf{C}}\big) = \operatorname{tr}(\mathbf{K}) - \operatorname{tr}\big(\bar{\mathbf{C}}\bar{\mathbf{C}}^\dagger\mathbf{K}\big) = \operatorname{tr}\big(\mathbf{K} - \mathbf{U}_{\bar{\mathbf{C}}}\mathbf{U}_{\bar{\mathbf{C}}}^T\mathbf{K}\big) \\ &= \operatorname{tr}\big[\mathbf{U}_{\bar{\mathbf{C}}}^\perp(\mathbf{U}_{\bar{\mathbf{C}}}^\perp)^T\mathbf{K}\big] = \operatorname{tr}\big[\mathbf{U}_{\bar{\mathbf{C}}}^\perp(\mathbf{U}_{\bar{\mathbf{C}}}^\perp)^T\mathbf{U}_{\bar{\mathbf{C}}}^\perp(\mathbf{U}_{\bar{\mathbf{C}}}^\perp)^T\mathbf{K}\big] = \operatorname{tr}\big[\mathbf{U}_{\bar{\mathbf{C}}}^\perp(\mathbf{U}_{\bar{\mathbf{C}}}^\perp)^T\mathbf{K}\mathbf{U}_{\bar{\mathbf{C}}}^\perp(\mathbf{U}_{\bar{\mathbf{C}}}^\perp)^T\big] \geq 0.\end{aligned}$$

To this end, we have shown SPSD.

Then we assume $\mathbf{K}$ is positive definite and prove that its approximation formed by SS is also positive definite. Since $\mathbf{U}_{\bar{\mathbf{C}}}^\perp(\mathbf{U}_{\bar{\mathbf{C}}}^\perp)^T\mathbf{K}\mathbf{U}_{\bar{\mathbf{C}}}^\perp(\mathbf{U}_{\bar{\mathbf{C}}}^\perp)^T$ is SPSD, its trace is zero only when it is all-zeros, which is equivalent to $\mathbf{K} = \mathbf{U}_{\bar{\mathbf{C}}}\mathbf{U}_{\bar{\mathbf{C}}}^T\mathbf{K}\mathbf{U}_{\bar{\mathbf{C}}}\mathbf{U}_{\bar{\mathbf{C}}}^T$. However, this cannot hold if $\mathbf{K}$ has full rank. Therefore $\delta^{\text{ss}} > 0$ holds when $\mathbf{K}$ is positive definite. When $\mathbf{K}$ is positive definite, the $c \times c$ matrix $\mathbf{U}_{\bar{\mathbf{C}}}^T\mathbf{K}\mathbf{U}_{\bar{\mathbf{C}}}$ is also positive definite. It follows from (29) that

$$\tilde{\mathbf{K}} = \bar{\mathbf{C}}\mathbf{U}^{\text{ss}}\bar{\mathbf{C}}^T + \delta^{\text{ss}}\mathbf{I}_n = \begin{bmatrix} \mathbf{U}_{\bar{\mathbf{C}}} & \mathbf{U}_{\bar{\mathbf{C}}}^\perp \end{bmatrix} \begin{bmatrix} \mathbf{U}_{\bar{\mathbf{C}}}^T\mathbf{K}\mathbf{U}_{\bar{\mathbf{C}}} & \mathbf{0} \\ \mathbf{0} & \delta^{\text{ss}}\mathbf{I}_{n-c} \end{bmatrix} \begin{bmatrix} \mathbf{U}_{\bar{\mathbf{C}}}^T \\ (\mathbf{U}_{\bar{\mathbf{C}}}^\perp)^T \end{bmatrix}.$$

Here the block diagonal matrix is positive definite, and thus $\tilde{\mathbf{K}}$ is positive definite.

Appendix F. Proof of Theorem 8 and Theorem 9

Directly analyzing the theoretical error bound of the SS model is not easy, so we formulate a variant of SS called the inexact spectral shifting (ISS) model and instead analyze the error bound of ISS. We define ISS in Section F.1 and prove Theorem 8 and Theorem 9 in Section F.2 and F.3, respectively. In the following we let $\tilde{\mathbf{K}}_c^{\text{ss}}$ and $\tilde{\mathbf{K}}_c^{\text{iss}}$ be respectively the approximation formed by the two models.

F.1 The ISS Model

ISS is defined by

$$\tilde{\mathbf{K}}_c^{\text{iss}} = \bar{\mathbf{C}}\bar{\mathbf{U}}\bar{\mathbf{C}}^T + \delta\mathbf{I}_n, \tag{30}$$

where $\delta > 0$ is the spectral shifting term, and $\bar{\mathbf{C}}\bar{\mathbf{U}}\bar{\mathbf{C}}^T$ is the prototype model of $\bar{\mathbf{K}} = \mathbf{K} - \delta\mathbf{I}_n$. It follows from Theorem 6 that ISS is less accurate than SS in that

$$\big\|\mathbf{K} - \tilde{\mathbf{K}}_c^{\text{ss}}\big\|_F \leq \big\|\mathbf{K} - \tilde{\mathbf{K}}_c^{\text{iss}}\big\|_F,$$

thus the error bounds of ISS still hold if $\tilde{\mathbf{K}}_c^{\text{iss}}$ is replaced by $\tilde{\mathbf{K}}_c^{\text{ss}}$.

We first show how to set the spectral shifting term δ. Since $\bar{\mathbf{C}}\bar{\mathbf{U}}\bar{\mathbf{C}}^T$ is the approximation to $\bar{\mathbf{K}}$ formed by the prototype model, it can holds with high probability that

$$\left\|\mathbf{K}-\tilde{\mathbf{K}}_c^{\mathrm{iss}}\right\|_F \;=\; \left\|\mathbf{K}-\bar{\delta}\mathbf{I}_n-\bar{\mathbf{C}}\bar{\mathbf{U}}\bar{\mathbf{C}}^T\right\|_F \;=\; \left\|\bar{\mathbf{K}}-\bar{\mathbf{C}}\bar{\mathbf{U}}\bar{\mathbf{C}}^T\right\|_F \;\leq\; \eta\left\|\bar{\mathbf{K}}-\bar{\mathbf{K}}_k\right\|_F$$

for some error parameter η. Apparently, for fixed k, the smaller the error $\|\bar{\mathbf{K}}-\bar{\mathbf{K}}_k\|_F$ is, the tighter error bound the ISS has; if $\|\bar{\mathbf{K}}-\bar{\mathbf{K}}_k\|_F \leq \|\mathbf{K}-\mathbf{K}_k\|_F$, then ISS has a better error bound than the prototype model. Therefore, our goal is to make $\|\bar{\mathbf{K}}-\bar{\mathbf{K}}_k\|_F$ as small as possible, so we formulate the following optimization problem to compute δ:

$$\min_{\delta\geq 0}\ \left\|\bar{\mathbf{K}}-\bar{\mathbf{K}}_k\right\|_F^2; \qquad \text{s.t. } \bar{\mathbf{K}}=\mathbf{K}-\delta\mathbf{I}_n.$$

However, since $\bar{\mathbf{K}}$ is in general indefinite, it requires all of the eigenvalues of $\mathbf{K}$ to solve the problem exactly. Since computing the full eigenvalue decomposition is expensive, we attempt to relax the problem. Considering that

$$\left\|\bar{\mathbf{K}}-\bar{\mathbf{K}}_k\right\|_F^2 \;=\; \min_{|\mathcal{J}|=n-k}\sum_{j\in\mathcal{J}}\big(\sigma_j(\mathbf{K})-\delta\big)^2 \;\leq\; \sum_{j=k+1}^{n}\big(\sigma_j(\mathbf{K})-\delta\big)^2, \tag{31}$$

we seek to minimize the upper bound of $\|\bar{\mathbf{K}}-\bar{\mathbf{K}}_k\|_F^2$, which is the right-hand side of (31), to compute δ, leading to the solution

$$\bar{\delta} \;=\; \frac{1}{n-k}\sum_{j=k+1}^{n}\sigma_j(\mathbf{K}) \;=\; \frac{1}{n-k}\left(\mathrm{tr}(\mathbf{K})-\sum_{j=1}^{k}\sigma_j(\mathbf{K})\right). \tag{32}$$

If we choose $\delta=0$, then ISS degenerates to the prototype model.

F.2 Proof of Theorem 8

Theorem 6 indicates that ISS is less accurate than SS, thus

$$\left\|\mathbf{K}-\tilde{\mathbf{K}}_c^{\mathrm{ss}}\right\|_F \;\leq\; \left\|\mathbf{K}-\tilde{\mathbf{K}}_c^{\mathrm{iss}}\right\|_F \;=\; \left\|\mathbf{K}-\bar{\delta}\mathbf{I}_n-\bar{\mathbf{C}}\bar{\mathbf{C}}^{\dagger}\bar{\mathbf{K}}(\bar{\mathbf{C}}^{\dagger})^T\bar{\mathbf{C}}^T\right\|_F \;=\; \left\|\bar{\mathbf{K}}-\bar{\mathbf{C}}\bar{\mathbf{C}}^{\dagger}\bar{\mathbf{K}}(\bar{\mathbf{C}}^{\dagger})^T\bar{\mathbf{C}}^T\right\|_F.$$

Theorem 8 follows from the above inequality and the following theorem.

Theorem 24 *Let $\mathbf{K}$ be any $n\times n$ SPSD matrix, $\bar{\delta}$ be defined in (32), and $\bar{\mathbf{K}}=\mathbf{K}-\delta\mathbf{I}_n$. Then for any $\delta\in(0,\bar{\delta}]$, the following inequality holds:*

$$\left\|\bar{\mathbf{K}}-\bar{\mathbf{K}}_k\right\|_F^2 \;\leq\; \left\|\mathbf{K}-\mathbf{K}_k\right\|_F^2.$$

Proof Since the righthand side of (31) is convex and $\bar{\delta}$ is the minimizer of the righthand of (31), for any $\delta\in(0,\bar{\delta}]$, it holds that

$$\sum_{j=k+1}^{n}\big(\sigma_j(\mathbf{K})-\delta\big)^2 \;\leq\; \sum_{j=k+1}^{n}\big(\sigma_j(\mathbf{K})-0\big)^2 \;=\; \left\|\mathbf{K}-\mathbf{K}_k\right\|_F^2.$$

Then the theorem follows by the inequality in (31). ■

F.3 Proof of Theorem 9

Since $\|\mathbf{K}-\tilde{\mathbf{K}}_c^{\mathrm{ss}}\|_F \le \|\mathbf{K}-\tilde{\mathbf{K}}_c^{\mathrm{iss}}\|_F$, Theorem 9 follows from Theorem 25.

Theorem 25 *Suppose there is a sketching matrix $\mathbf{P}\in\mathbb{R}^{n\times c}$ such that for any $n\times n$ symmetric matrix $\mathbf{S}$ and target rank k ($\ll n$), by forming the sketch $\mathbf{C}=\mathbf{SP}$, the prototype model satisfies the error bound*

$$\left\|\mathbf{S}-\mathbf{C}\mathbf{C}^{\dagger}\mathbf{S}(\mathbf{C}^{\dagger})^T\mathbf{C}^T\right\|_F^2 \le \eta\left\|\mathbf{S}-\mathbf{S}_k\right\|_F^2.$$

Let $\mathbf{K}$ be any $n\times n$ SPSD matrix, $\bar{\delta}$ defined in (32) be the initial spectral shifting term, and $\tilde{\mathbf{K}}_c^{iss}$ be the ISS approximation defined in (30). Then

$$\left\|\mathbf{K}-\tilde{\mathbf{K}}_c^{iss}\right\|_F^2 \le \eta\left(\left\|\mathbf{K}-\mathbf{K}_k\right\|_F^2-\frac{\left[\sum_{i=k+1}^n\lambda_i(\mathbf{K})\right]^2}{n-k}\right).$$

Proof The error incurred by SS is

$$\begin{aligned}\left\|\mathbf{K}-\tilde{\mathbf{K}}_c^{\mathrm{ss}}\right\|_F^2 &= \left\|(\bar{\mathbf{K}}+\bar{\delta}\mathbf{I}_n)-(\bar{\mathbf{C}}\bar{\mathbf{U}}\bar{\mathbf{C}}^T+\bar{\delta}\mathbf{I}_n)\right\|_F^2 = \left\|\bar{\mathbf{K}}-\bar{\mathbf{C}}\bar{\mathbf{U}}\bar{\mathbf{C}}^T\right\|_F^2\\ &\le \eta\left\|\bar{\mathbf{K}}-\bar{\mathbf{K}}_k\right\|_F^2 = \eta\sum_{i=k+1}^n\sigma_i^2(\bar{\mathbf{K}}) = \eta\sum_{i=k+1}^n\lambda_i(\bar{\mathbf{K}}^2).\end{aligned}$$

Here the inequality follows from the assumption. The i-th largest eigenvalue of $\bar{\mathbf{K}}$ is $\lambda_i(\mathbf{K})-\bar{\delta}$, so the n eigenvalues of $\bar{\mathbf{K}}^2$ are all in the set $\{(\lambda_i(\mathbf{K})-\bar{\delta})^2\}_{i=1}^n$. The sum of the smallest $n-k$ of the n eigenvalues of $\bar{\mathbf{K}}^2$ must be less than or equal to the sum of any $n-k$ of the eigenvalues, thus we have

$$\begin{aligned}\sum_{i=k+1}^n\lambda_i(\bar{\mathbf{K}}^2) &\le \sum_{i=k+1}^n\left(\lambda_i(\mathbf{K})-\bar{\delta}\right)^2\\ &= \sum_{i=k+1}^n\lambda_i^2(\mathbf{K})-2\sum_{i=k+1}^n\bar{\delta}\lambda_i(\mathbf{K})+(n-k)(\bar{\delta})^2\\ &= \left\|\mathbf{K}-\mathbf{K}_k\right\|_F^2-\frac{1}{n-k}\left[\sum_{i=k+1}^n\lambda_i(\mathbf{K})\right]^2,\end{aligned}$$

by which the theorem follows. ■

Appendix G. Proof of Theorem 11

Proof Let $\tilde{\mathbf{K}}=\mathbf{Q}(\mathbf{Q}^T\mathbf{K})_k$, where $\mathbf{Q}$ is defined in Line 4 in Algorithm 5. Boutsidis et al. (2014) showed that

$$\mathbb{E}\|\mathbf{K}-\tilde{\mathbf{K}}\|_F^2 \le (1+k/l)\left\|\mathbf{K}-\mathbf{K}_k\right\|_F^2, \tag{33}$$

where the expectation is taken w.r.t. the random Gaussian matrix $\mathbf{\Omega}$.

It follows from Lemma 26 that

$$\|\boldsymbol{\sigma}_{\mathbf{K}} - \boldsymbol{\sigma}_{\tilde{\mathbf{K}}}\|_2^2 \;\leq\; \|\mathbf{K} - \tilde{\mathbf{K}}\|_F^2,$$

where $\boldsymbol{\sigma}_{\mathbf{K}}$ and $\boldsymbol{\sigma}_{\tilde{\mathbf{K}}}$ contain the singular values in a descending order. Since $\tilde{\mathbf{K}}$ has a rank at most k, the $k+1$ to n entries of $\boldsymbol{\sigma}_{\tilde{\mathbf{K}}}$ are zero. We split $\boldsymbol{\sigma}_{\mathbf{K}}$ and $\boldsymbol{\sigma}_{\tilde{\mathbf{K}}}$ into vectors of length k and $n-k$:

$$\boldsymbol{\sigma}_{\mathbf{K}} \;=\; \begin{bmatrix} \boldsymbol{\sigma}_{\mathbf{K},k} \\ \boldsymbol{\sigma}_{\mathbf{K},-k} \end{bmatrix} \quad \text{and} \quad \boldsymbol{\sigma}_{\tilde{\mathbf{K}}} \;=\; \begin{bmatrix} \boldsymbol{\sigma}_{\tilde{\mathbf{K}},k} \\ \mathbf{0} \end{bmatrix}$$

and thus

$$\|\boldsymbol{\sigma}_{\mathbf{K},k} - \boldsymbol{\sigma}_{\tilde{\mathbf{K}},k}\|_2^2 + \|\boldsymbol{\sigma}_{\mathbf{K},-k}\|_2^2 \;\leq\; \|\mathbf{K} - \tilde{\mathbf{K}}\|_F^2. \tag{34}$$

Since $\|\boldsymbol{\sigma}_{\mathbf{K},-k}\|_2^2 = \|\mathbf{K} - \mathbf{K}_k\|_F^2$, it follows from (33) and (34) that

$$\mathbb{E}\|\boldsymbol{\sigma}_{\mathbf{K},k} - \boldsymbol{\sigma}_{\tilde{\mathbf{K}},k}\|_2^2 \;\leq\; \frac{k}{l}\,\|\boldsymbol{\sigma}_{\mathbf{K},-k}\|_2^2.$$

Since $\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_1 \leq \sqrt{k}\|\mathbf{x}\|_2$ for any $\mathbf{x} \in \mathbb{R}^k$, we have that

$$\mathbb{E}\big\|\boldsymbol{\sigma}_{\mathbf{K},k} - \boldsymbol{\sigma}_{\tilde{\mathbf{K}},k}\big\|_1 \;\leq\; \frac{k}{\sqrt{l}}\,\big\|\boldsymbol{\sigma}_{\mathbf{K},-k}\big\|_1.$$

Then it follows from (32) and Line 6 in Algorithm 5 that

$$\begin{aligned}
\mathbb{E}\big|\bar{\delta} - \tilde{\delta}\big| \;&=\; \mathbb{E}\left[\frac{1}{n-k}\left|\sum_{i=1}^{k}\sigma_i(\mathbf{K}) - \sum_{i=1}^{k}\sigma_i(\tilde{\mathbf{K}})\right|\right] \\
&\leq\; \frac{1}{n-k}\,\mathbb{E}\big\|\boldsymbol{\sigma}_{\mathbf{K},k} - \boldsymbol{\sigma}_{\tilde{\mathbf{K}},k}\big\|_1 \;\leq\; \frac{k}{\sqrt{l}}\,\frac{1}{n-k}\,\big\|\boldsymbol{\sigma}_{\mathbf{K},-k}\big\|_1 \;=\; \frac{k}{\sqrt{l}}\,\bar{\delta}.
\end{aligned}$$

■

Lemma 26 *Let* $\mathbf{A}$ *and* $\mathbf{B}$ *be* $n \times n$ *matrices and* $\boldsymbol{\sigma}_{\mathbf{A}}$ *and* $\boldsymbol{\sigma}_{\mathbf{B}}$ *contain the singular values in a descending order. Then we have that*

$$\|\boldsymbol{\sigma}_{\mathbf{A}} - \boldsymbol{\sigma}_{\mathbf{B}}\|_2^2 \;\leq\; \|\mathbf{A} - \mathbf{B}\|_F^2.$$

Proof It is easy to show that

$$\begin{aligned}
\|\mathbf{A} - \mathbf{B}\|_F^2 \;&=\; \mathrm{tr}(\mathbf{A}^T\mathbf{A}) + \mathrm{tr}(\mathbf{B}^T\mathbf{B}) - 2\mathrm{tr}(\mathbf{A}^T\mathbf{B}) \\
&=\; \sum_{i=1}^{n}\sigma_i^2(\mathbf{A}) + \sum_{i=1}^{n}\sigma_i^2(\mathbf{B}) - 2\mathrm{tr}(\mathbf{A}^T\mathbf{B}).
\end{aligned} \tag{35}$$

We also have

$$\mathrm{tr}(\mathbf{A}^T\mathbf{B}) \;\leq\; \sum_{i=1}^{n}\sigma_i(\mathbf{A}^T\mathbf{B}) \;\leq\; \sum_{i=1}^{n}\sigma_i(\mathbf{A})\sigma_i(\mathbf{B}), \tag{36}$$

where the first inequality follows from Theorem 3.3.13 of Horn and Johnson and the second inequality follows from Theorem 3.3.14 of Horn and Johnson. Combining (35) and (36) we have that

$$\|\mathbf{A}-\mathbf{B}\|_F^2 \;\geq\; \sum_{i=1}^{n}\Big(\sigma_i^2(\mathbf{A})+\sigma_i^2(\mathbf{B})-2\sigma_i(\mathbf{A})\sigma_i(\mathbf{B})\Big) \;=\; \|\boldsymbol{\sigma}_{\mathbf{A}}-\boldsymbol{\sigma}_{\mathbf{B}}\|_2^2,$$

by which the theorem follows. ■

References

Adi Ben-Israel and Thomas N.E. Greville. *Generalized Inverses: Theory and Applications. Second Edition*. Springer, 2003. 31

Christos Boutsidis and David P. Woodruff. Optimal CUR matrix decompositions. *arXiv preprint arXiv:1405.7910*, 2014. 10, 34

Christos Boutsidis, Petros Drineas, and Malik Magdon-Ismail. Near-optimal column-based matrix reconstruction. *SIAM Journal on Computing*, 43(2):687–717, 2014. 9, 10, 34, 39, 45

Amit Deshpande, Luis Rademacher, Santosh Vempala, and Grant Wang. Matrix approximation and projective clustering via volume sampling. *Theory of Computing*, 2(2006): 225–247, 2006. 9, 40

Petros Drineas and Michael W. Mahoney. On the Nyström method for approximating a gram matrix for improved kernel-based learning. *Journal of Machine Learning Research*, 6:2153–2175, 2005. 2, 7

Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Relative-error CUR matrix decompositions. *SIAM Journal on Matrix Analysis and Applications*, 30(2):844–881, September 2008. 38

Petros Drineas, Malik Magdon-Ismail, Michael W. Mahoney, and David P. Woodruff. Fast approximation of matrix coherence and statistical leverage. *Journal of Machine Learning Research*, 13:3441–3472, 2012. 11

Alex Gittens and Michael W. Mahoney. Revisiting the nyström method for improved large-scale machine learning. In *International Conference on Machine Learning (ICML)*, 2013. 2, 7, 14

Ming Gu and Stanley C. Eisenstat. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM Journal on Scientific Computing*, 17(4):848–869, 1996. 9

Venkatesan Guruswami and Ali Kemal Sinop. Optimal column-based low-rank matrix reconstruction. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2012. 9

Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011. 2, 7

Roger A. Horn and Charles R. Johnson. Topics in matrix analysis. 1991. *Cambridge University Presss, Cambridge*. 47

Rong Jin, Tianbao Yang, Mehrdad Mahdavi, Y Li, and Z Zhou. Improved bounds for the Nyström method with application to kernel classification. *IEEE Transactions on Information Theory*, 59(10):6939–6949, 2013. 7

Sanjiv Kumar, Mehryar Mohri, and Ameet Talwalkar. On sampling-based approximate spectral decomposition. In *International Conference on Machine Learning (ICML)*, 2009. 3, 8

Sanjiv Kumar, Mehryar Mohri, and Ameet Talwalkar. Sampling methods for the Nyström method. *Journal of Machine Learning Research*, 13:981–1006, 2012. 4, 7, 9, 10, 12, 22, 23

David Lopez-Paz, Suvrit Sra, Alex Smola, Zoubin Ghahramani, and Bernhard Schölkopf. Randomized nonlinear component analysis. In *International Conference on Machine Learning (ICML)*, 2014. 2

Michael W. Mahoney. Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning*, 3(2):123–224, 2011. 9

Evert J. Nyström. Über die praktische auflösung von integralgleichungen mit anwendungen auf randwertaufgaben. *Acta Mathematica*, 54(1):185–204, 1930. 2, 6

Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: replacing minimization with randomization in learning. In *Advances in Neural Information Processing Systems (NIPS)*. 2008. 2

Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, 2002. 1

John Shawe-Taylor and Nello Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004. 1

John Shawe-Taylor, Christopher K. I. Williams, Nello Cristianini, and Jaz Kandola. On the eigenspectrum of the gram matrix and the generalisation error of kernel pca. *IEEE Transactions on Information Theory*, 51:2510–2522, 2005. 7

Si Si, Cho-Jui Hsieh, and Inderjit Dhillon. Memory efficient kernel approximation. In *International Conference on Machine Learning (ICML)*, pages 701–709, 2014. 4, 22, 23

G. W. Stewart. Four algorithms for the efficient computation of truncated pivoted QR approximations to a sparse matrix. *Numerische Mathematik*, 83(2):313–323, 1999. 9

Ameet Talwalkar and Afshin Rostamizadeh. Matrix coherence and the Nyström method. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2010. 3, 8

Ameet Talwalkar, Sanjiv Kumar, Mehryar Mohri, and Henry Rowley. Large-scale SVD and manifold learning. *Journal of Machine Learning Research*, 14:3129–3152, 2013. 7

Joel A Tropp. An introduction to matrix concentration inequalities. *arXiv preprint arXiv:1501.01571*, 2015. 36

Shusen Wang and Zhihua Zhang. Improving CUR matrix decomposition and the Nyström approximation via adaptive sampling. *Journal of Machine Learning Research*, 14:2729–2769, 2013. 2, 3, 4, 7, 9, 31, 35, 41

Shusen Wang, Zhihua Zhang, and Tong Zhang. Towards more efficient SPSD matrix approximation and CUR matrix decomposition. *arXiv preprint arXiv:1503.08395*, 2015. 3, 7, 8, 29

Christopher Williams and Matthias Seeger. Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems (NIPS)*, 2001. 2, 6

David P Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in Theoretical Computer Science*, 10(1–2):1–157, 2014. 9, 36

Tianbao Yang, Yu-Feng Li, Mehrdad Mahdavi, Rong Jin, and Zhi-Hua Zhou. Nyström method vs random fourier features: A theoretical and empirical comparison. In *Advances in Neural Information Processing Systems (NIPS)*, 2012. 2

Zhihua Zhang. The matrix ridge approximation: algorithms and applications. *Machine Learning*, 97:227–258, 2014. 4